

US 20230134266A1

(19) **United States**

(12) **Patent Application Publication**
Sundararajan et al.

(10) **Pub. No.: US 2023/0134266 A1**

(43) **Pub. Date: May 4, 2023**

(54) **STORAGE OPERATION ROUTING IN A CONTAINER SYSTEM**

(71) Applicant: **Pure Storage, Inc.**, Mountain View, CA (US)

(72) Inventors: **Lakshmi Narasimhan Sundararajan**, Chennai (IN); **Prabir Paul**, San Jose, CA (US); **Nikhil Subhash Bhupale**, Pune (IN)

(21) Appl. No.: **17/853,466**

(22) Filed: **Jun. 29, 2022**

Related U.S. Application Data

(63) Continuation-in-part of application No. 17/514,872, filed on Oct. 29, 2021.

Publication Classification

(51) **Int. Cl.**
G06F 3/06 (2006.01)

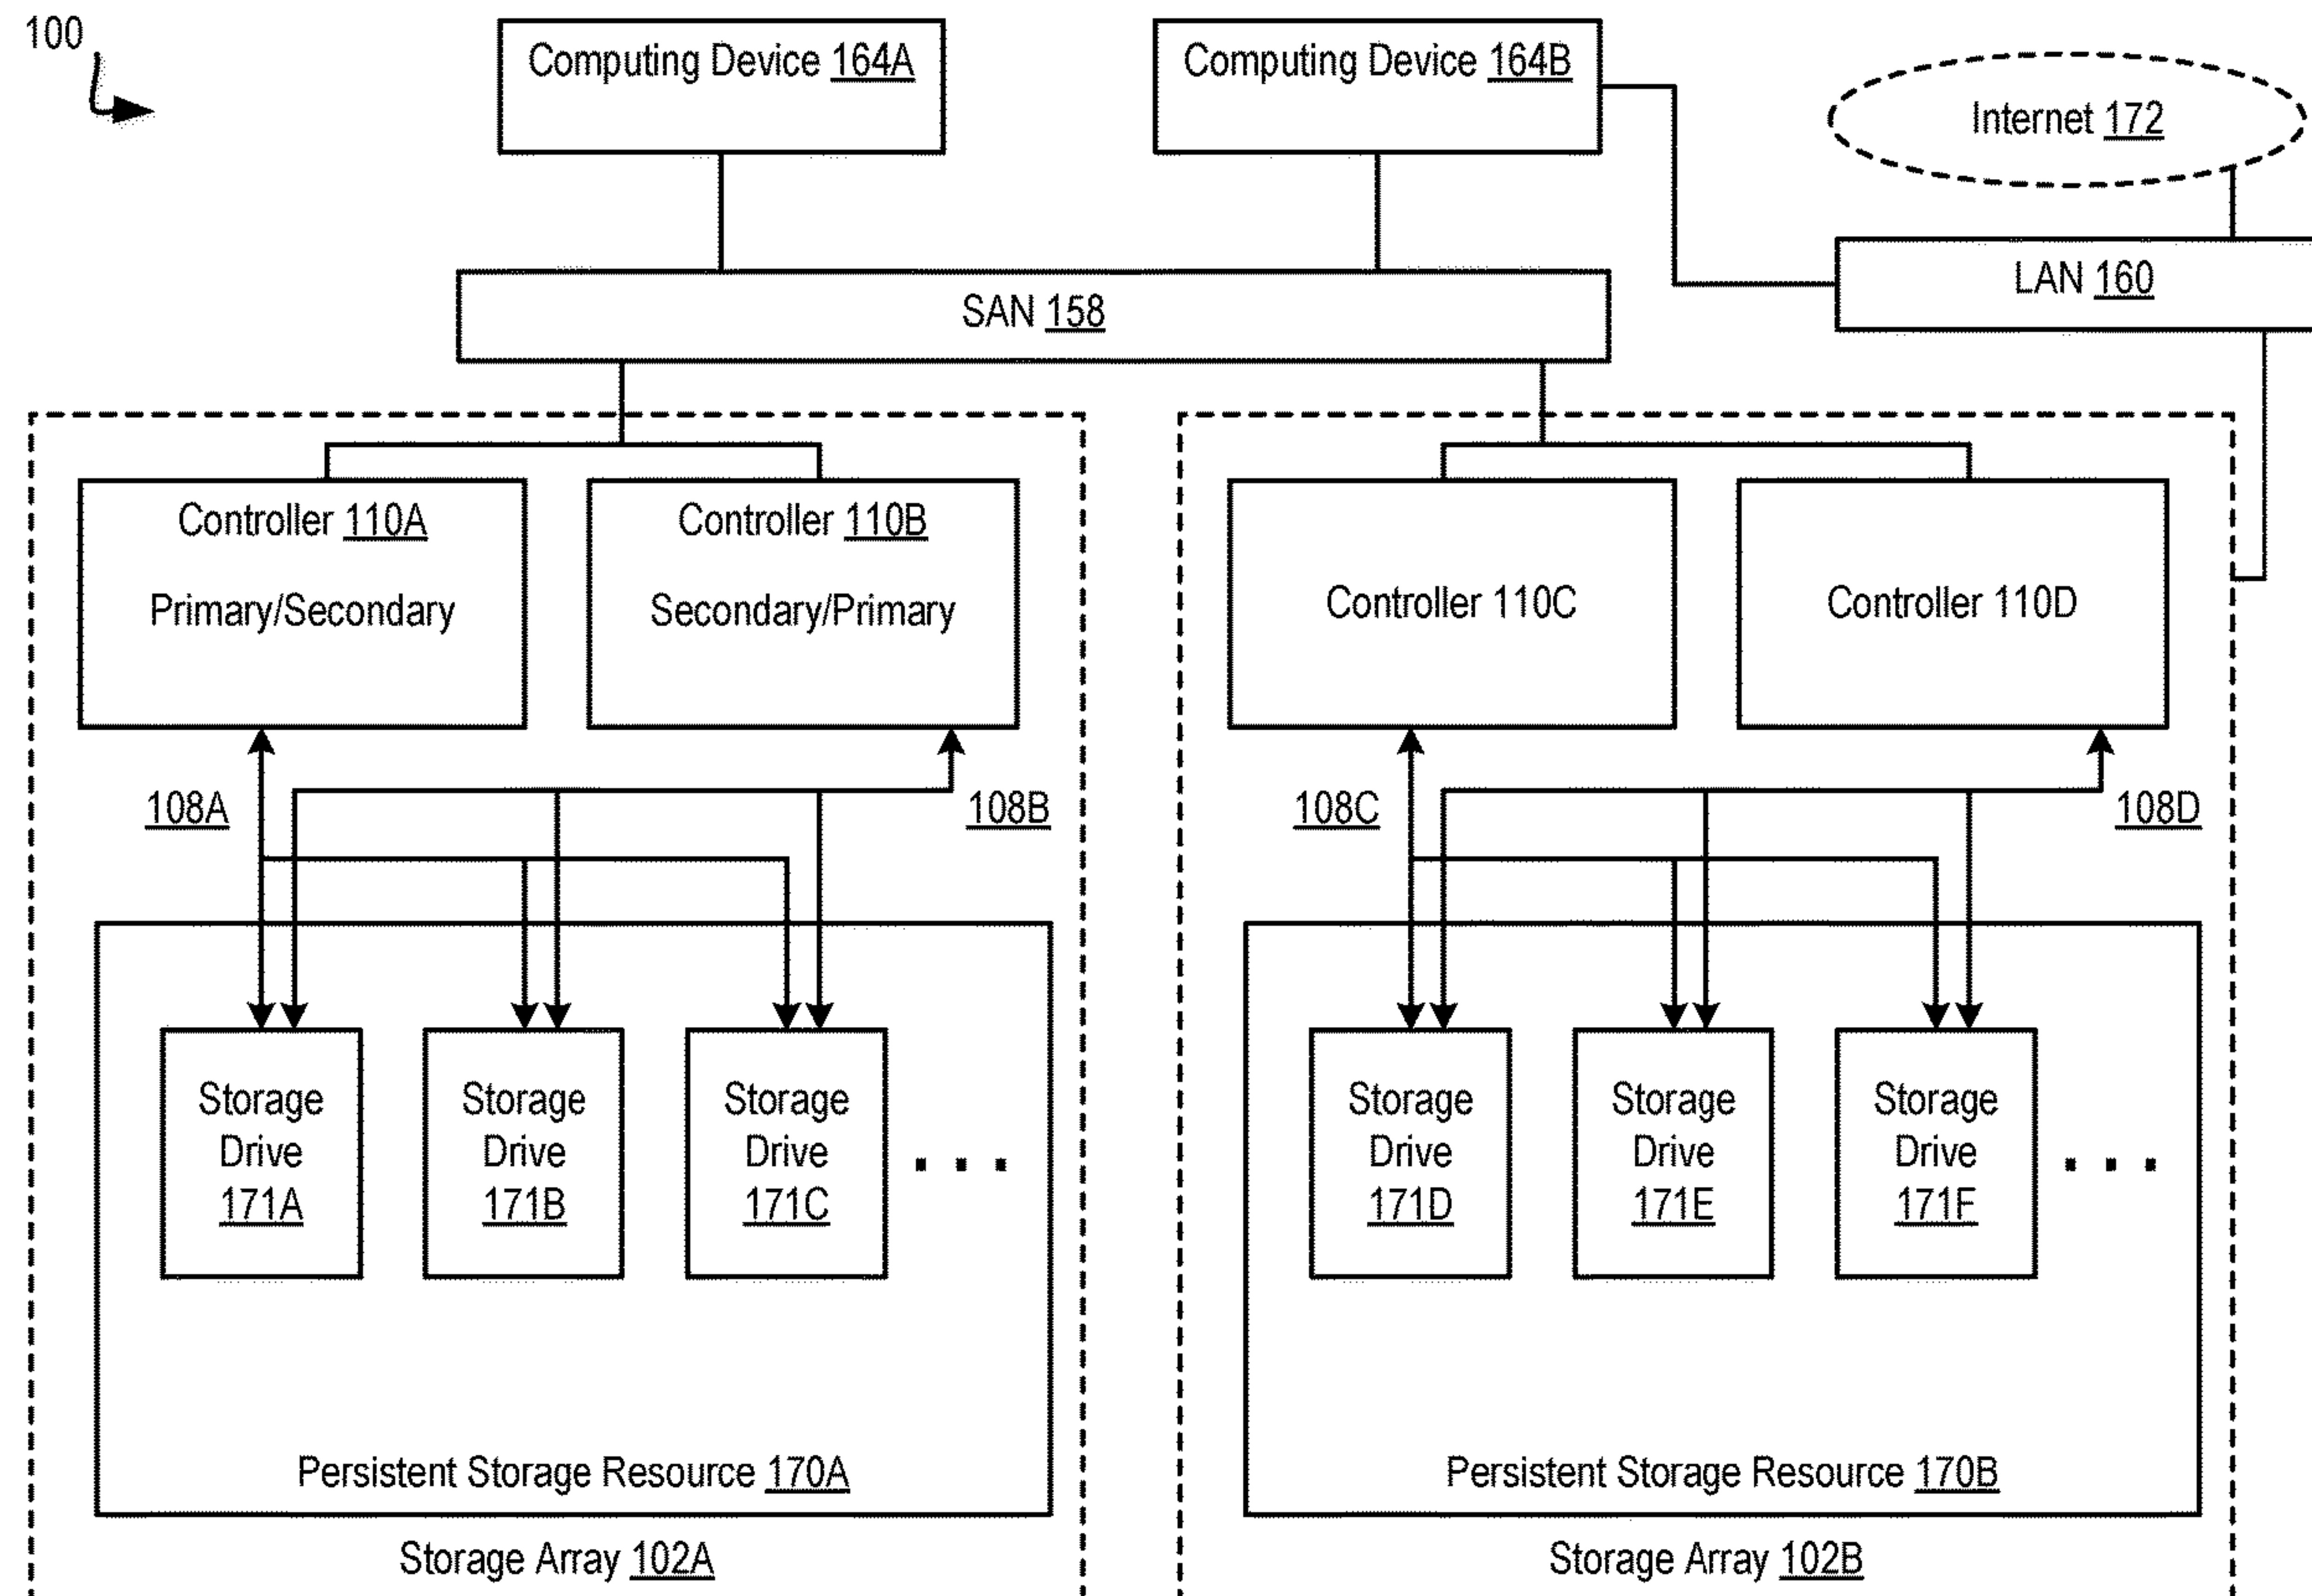
(52) **U.S. Cl.**

CPC **G06F 3/0635** (2013.01); **G06F 3/0643** (2013.01); **G06F 3/065** (2013.01); **G06F 3/0659** (2013.01); **G06F 3/0604** (2013.01); **G06F 3/067** (2013.01)

(57)

ABSTRACT

In some embodiments, storage operation routing in a container system may include receiving, by a container storage system executing on a node within a container system, a storage operation configured to be performed with respect to a storage resource; determining whether the storage resource is locally accessible at the node; and routing the storage operation based on the determining of whether the storage resource is locally accessible at the node. For example, the routing may be performed such that when the storage resource is not locally accessible at the node, the routing is by way of both a kernel space and a user space of the node, and when the storage resource is locally accessible at the node, the routing is contained within the kernel space of the node. Corresponding methods and systems are also disclosed.



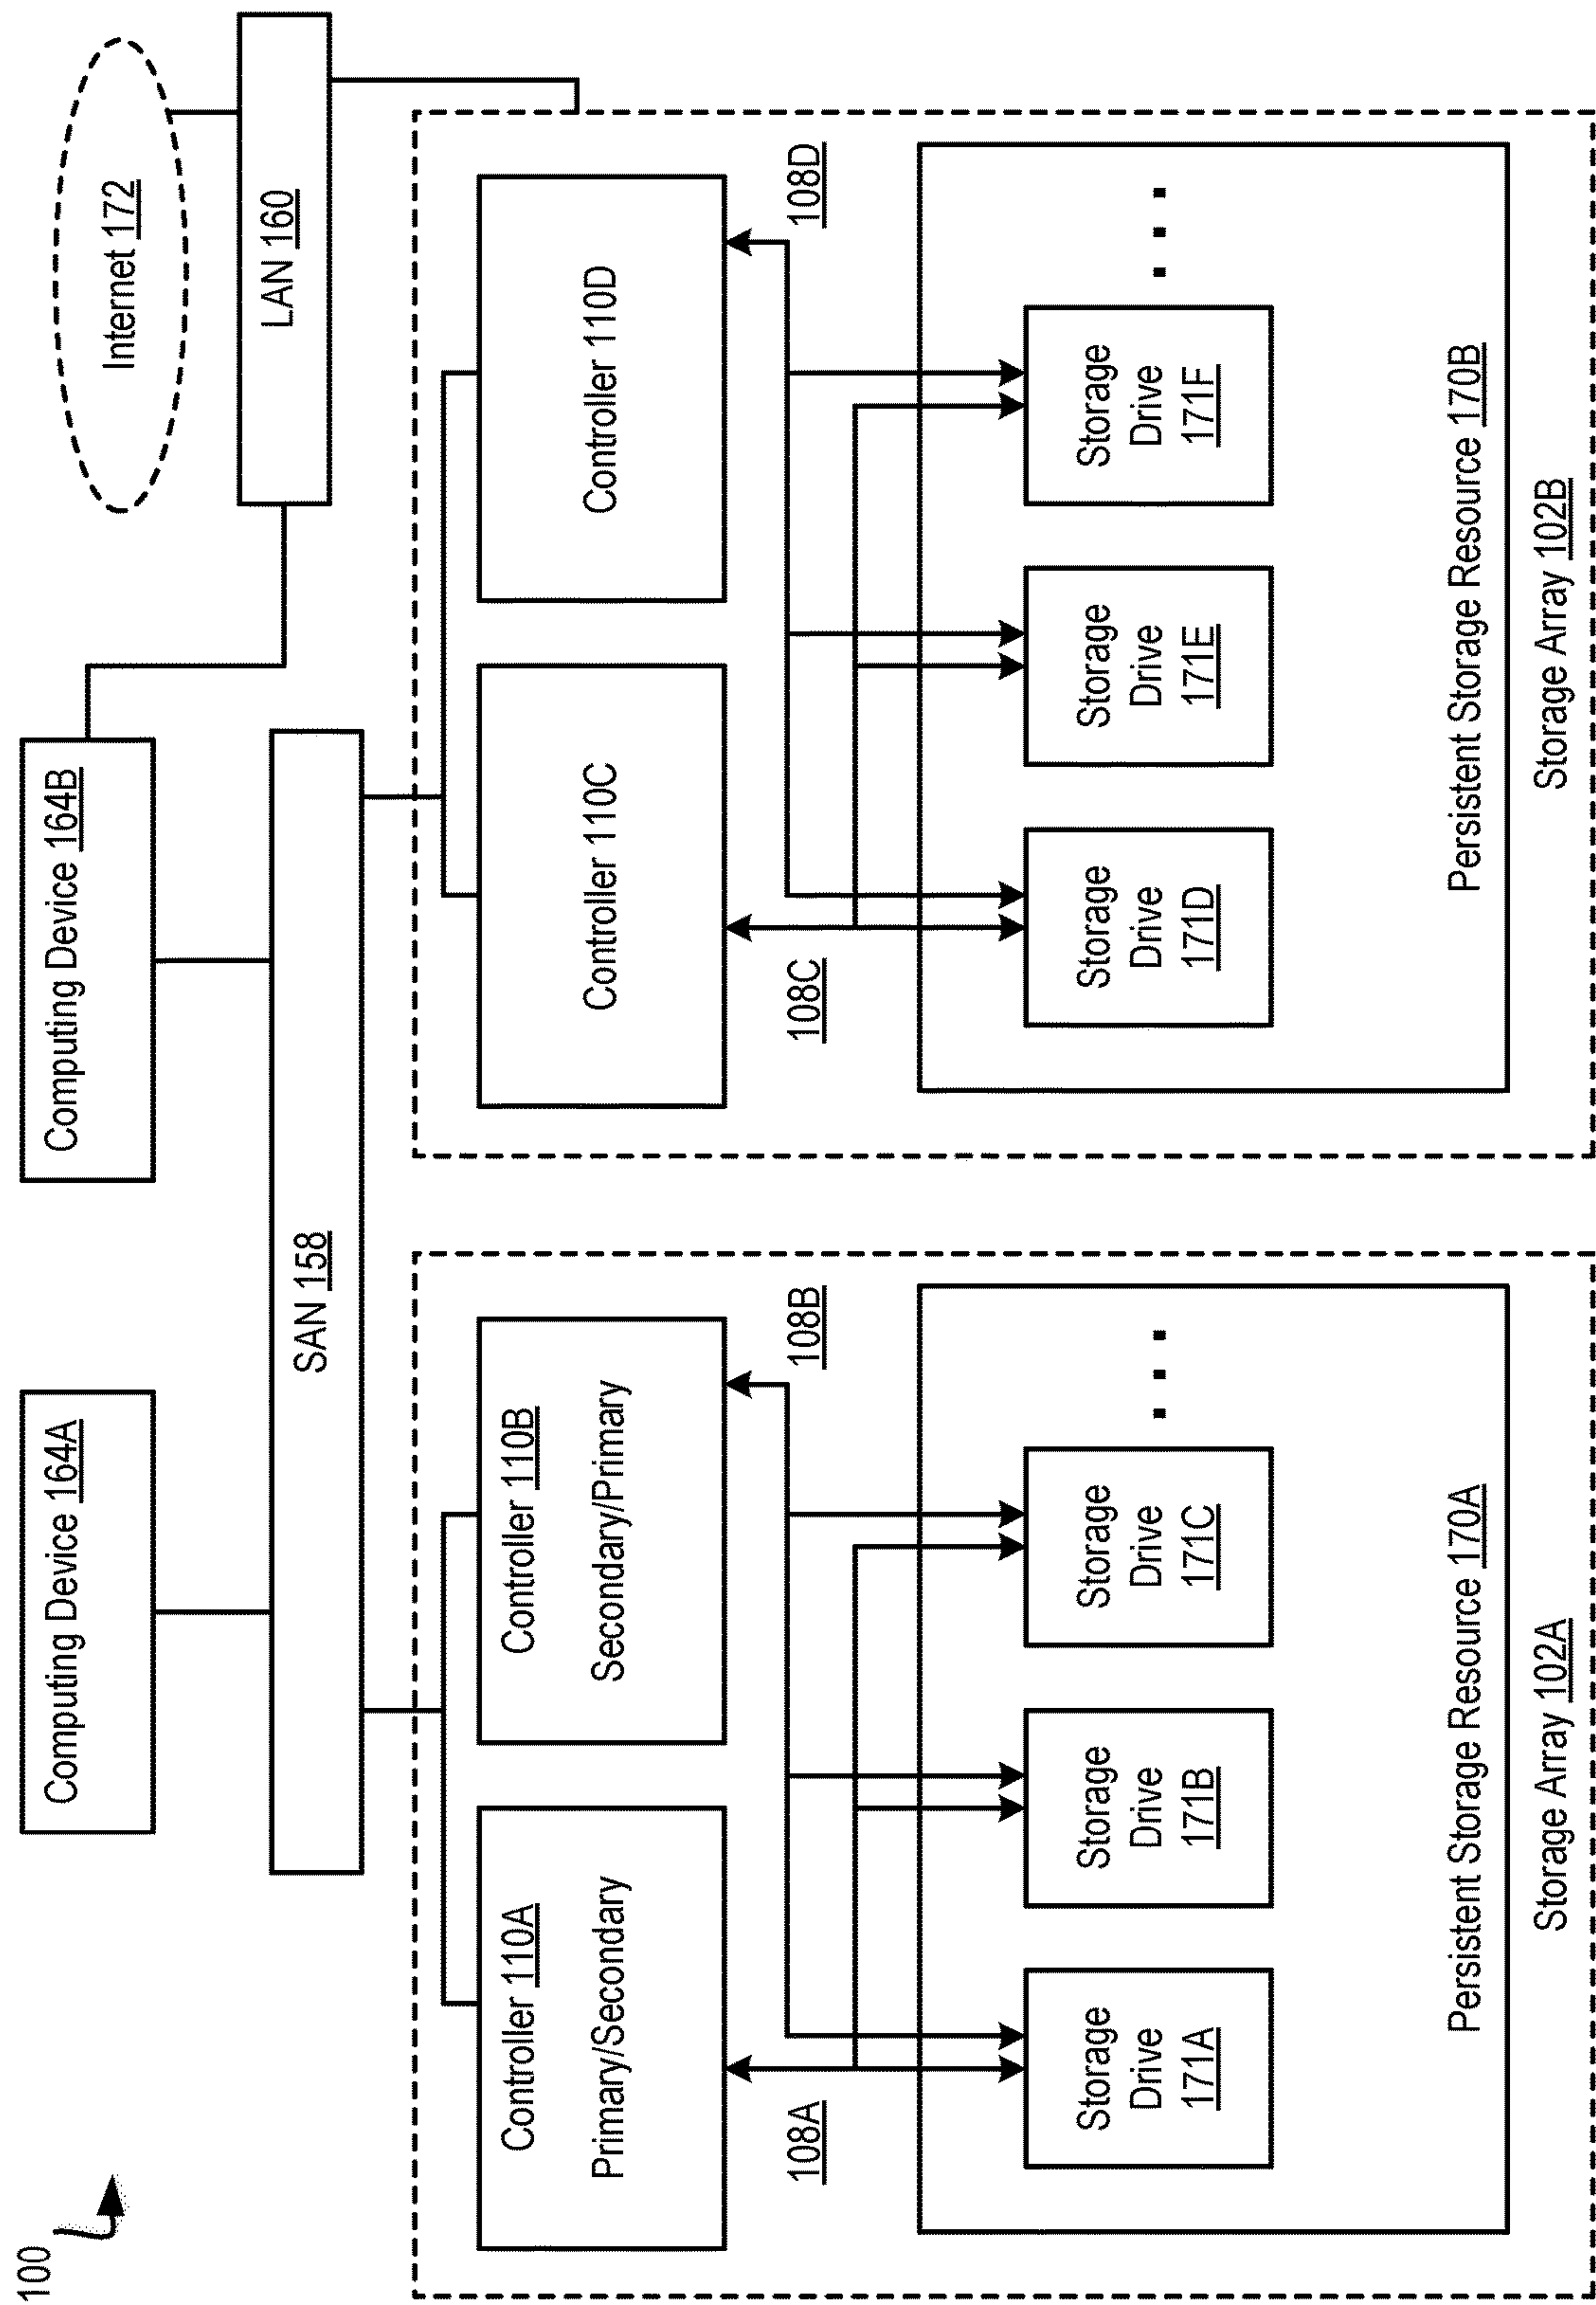


FIG. 1A




FIG. 1B




FIG. 1C

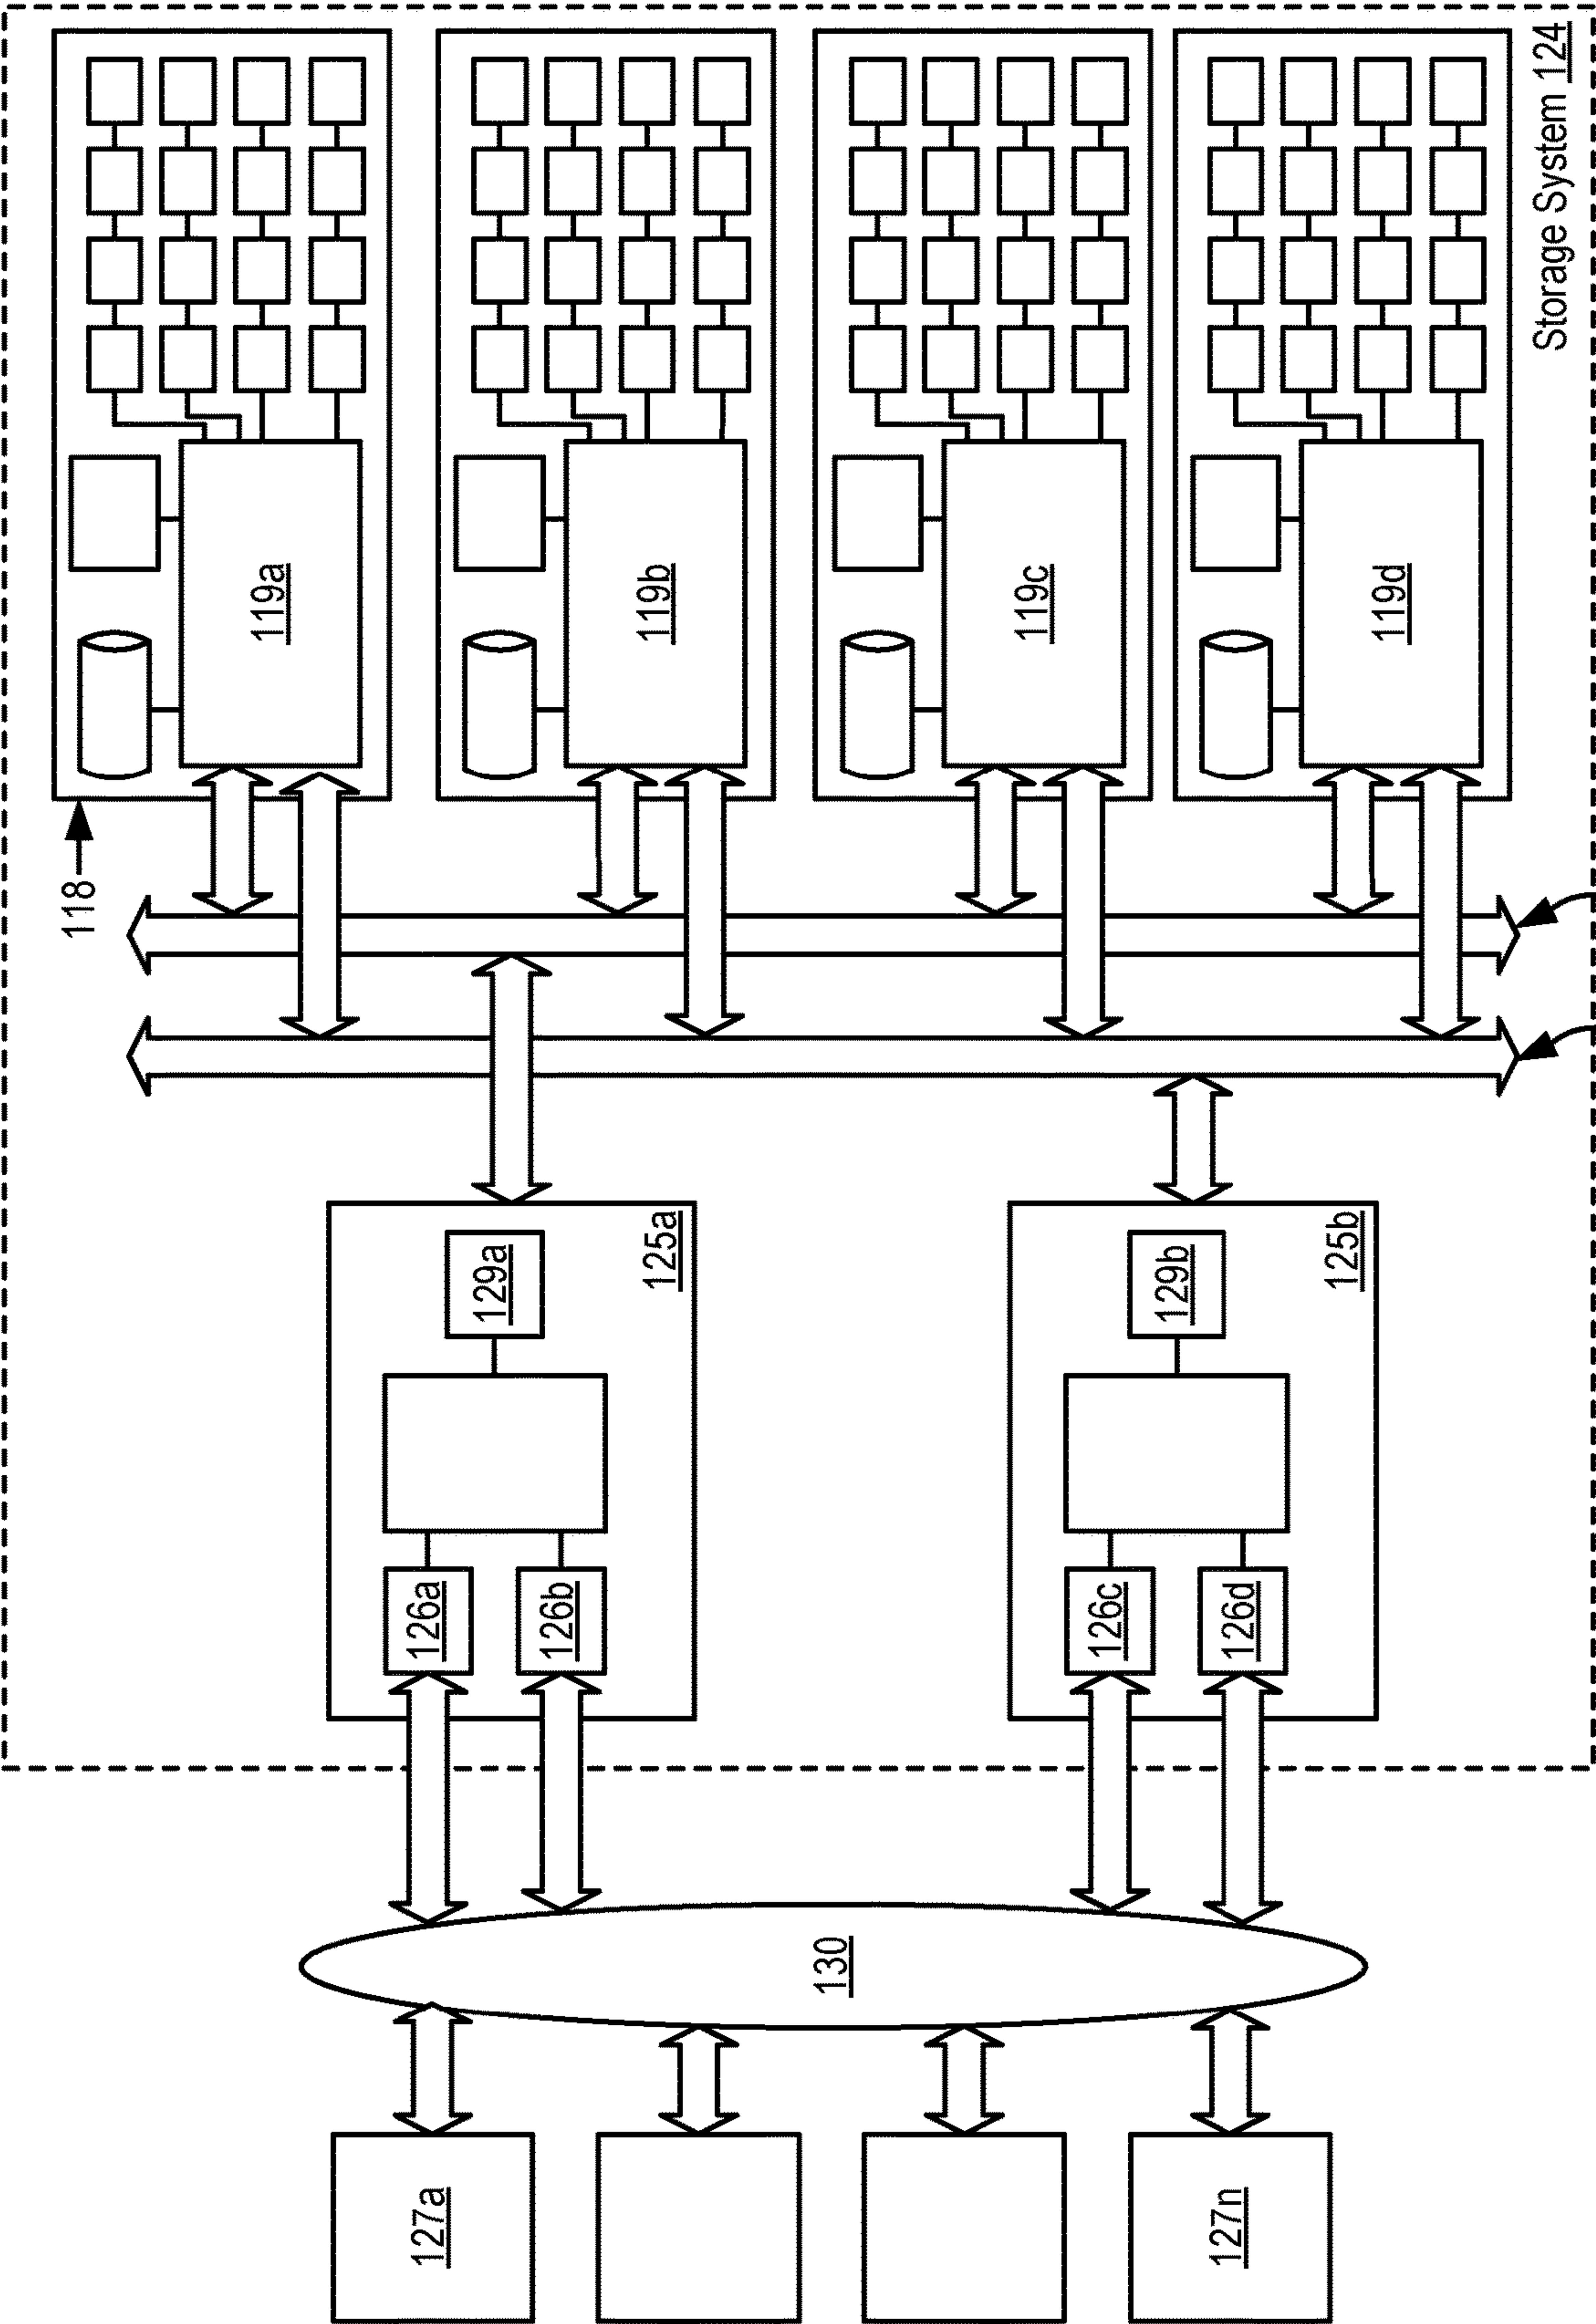


FIG. 1D




FIG. 2A




FIG. 2B




FIG. 2C

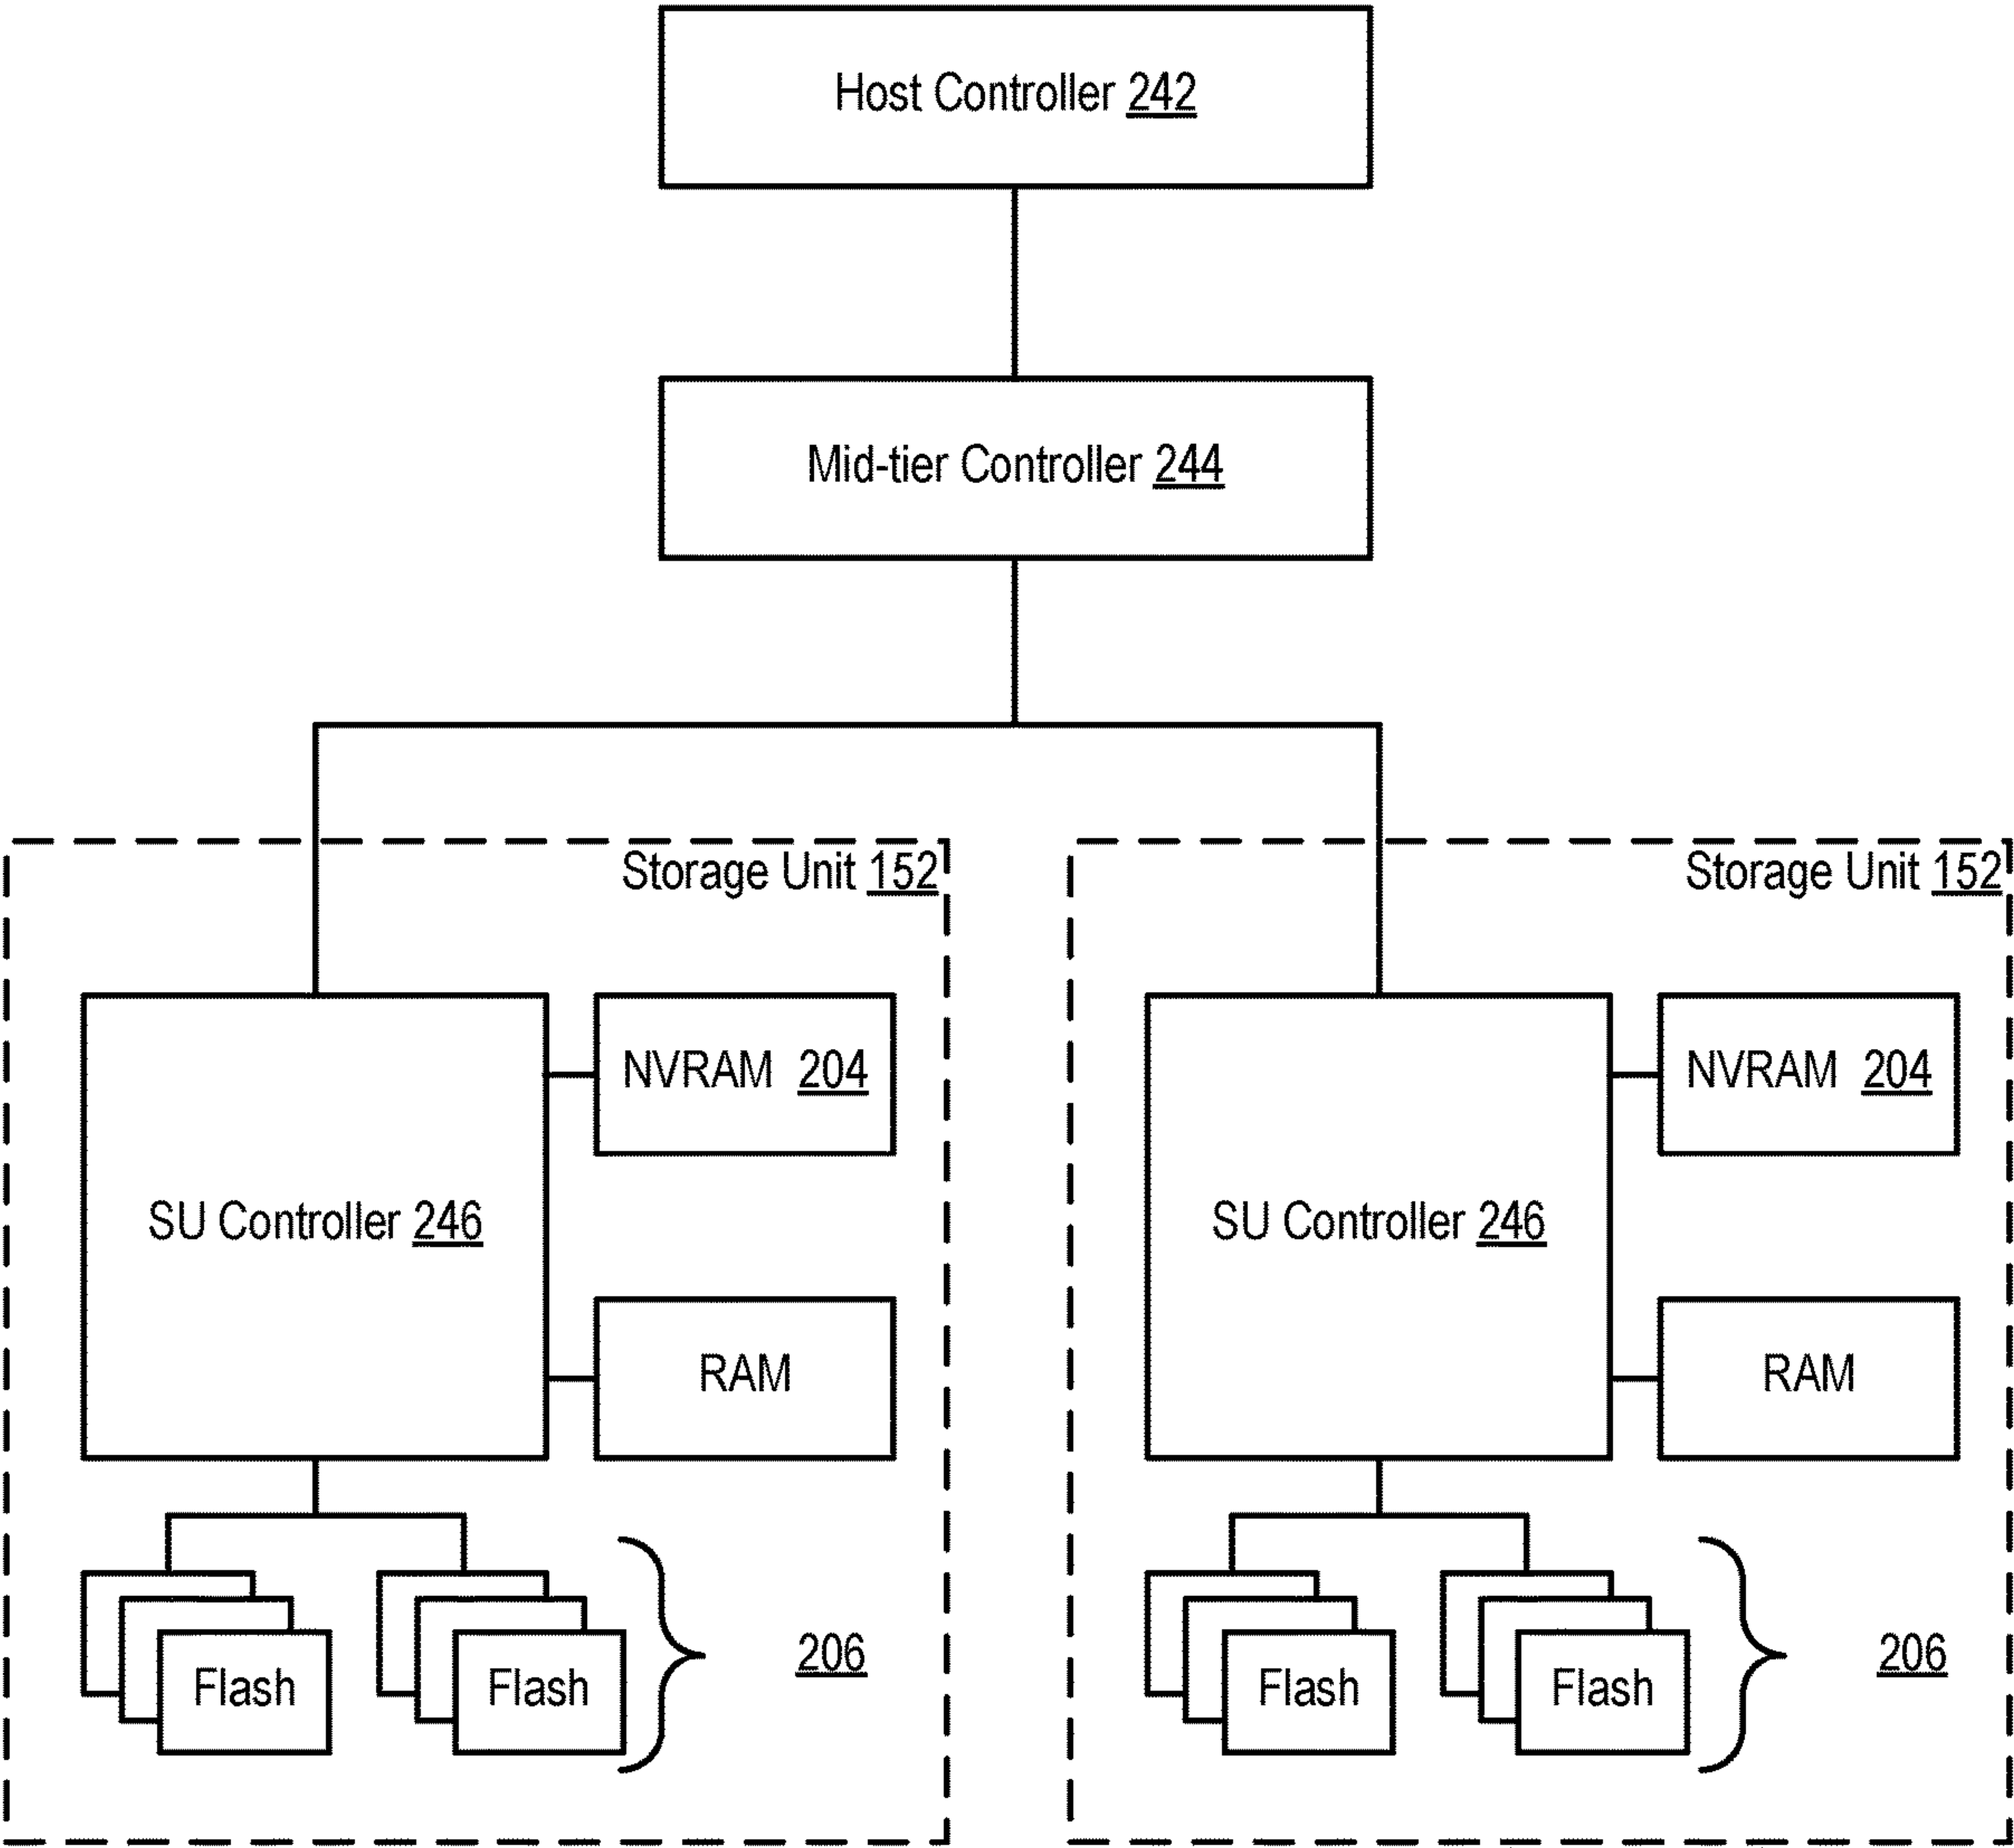



FIG. 2D



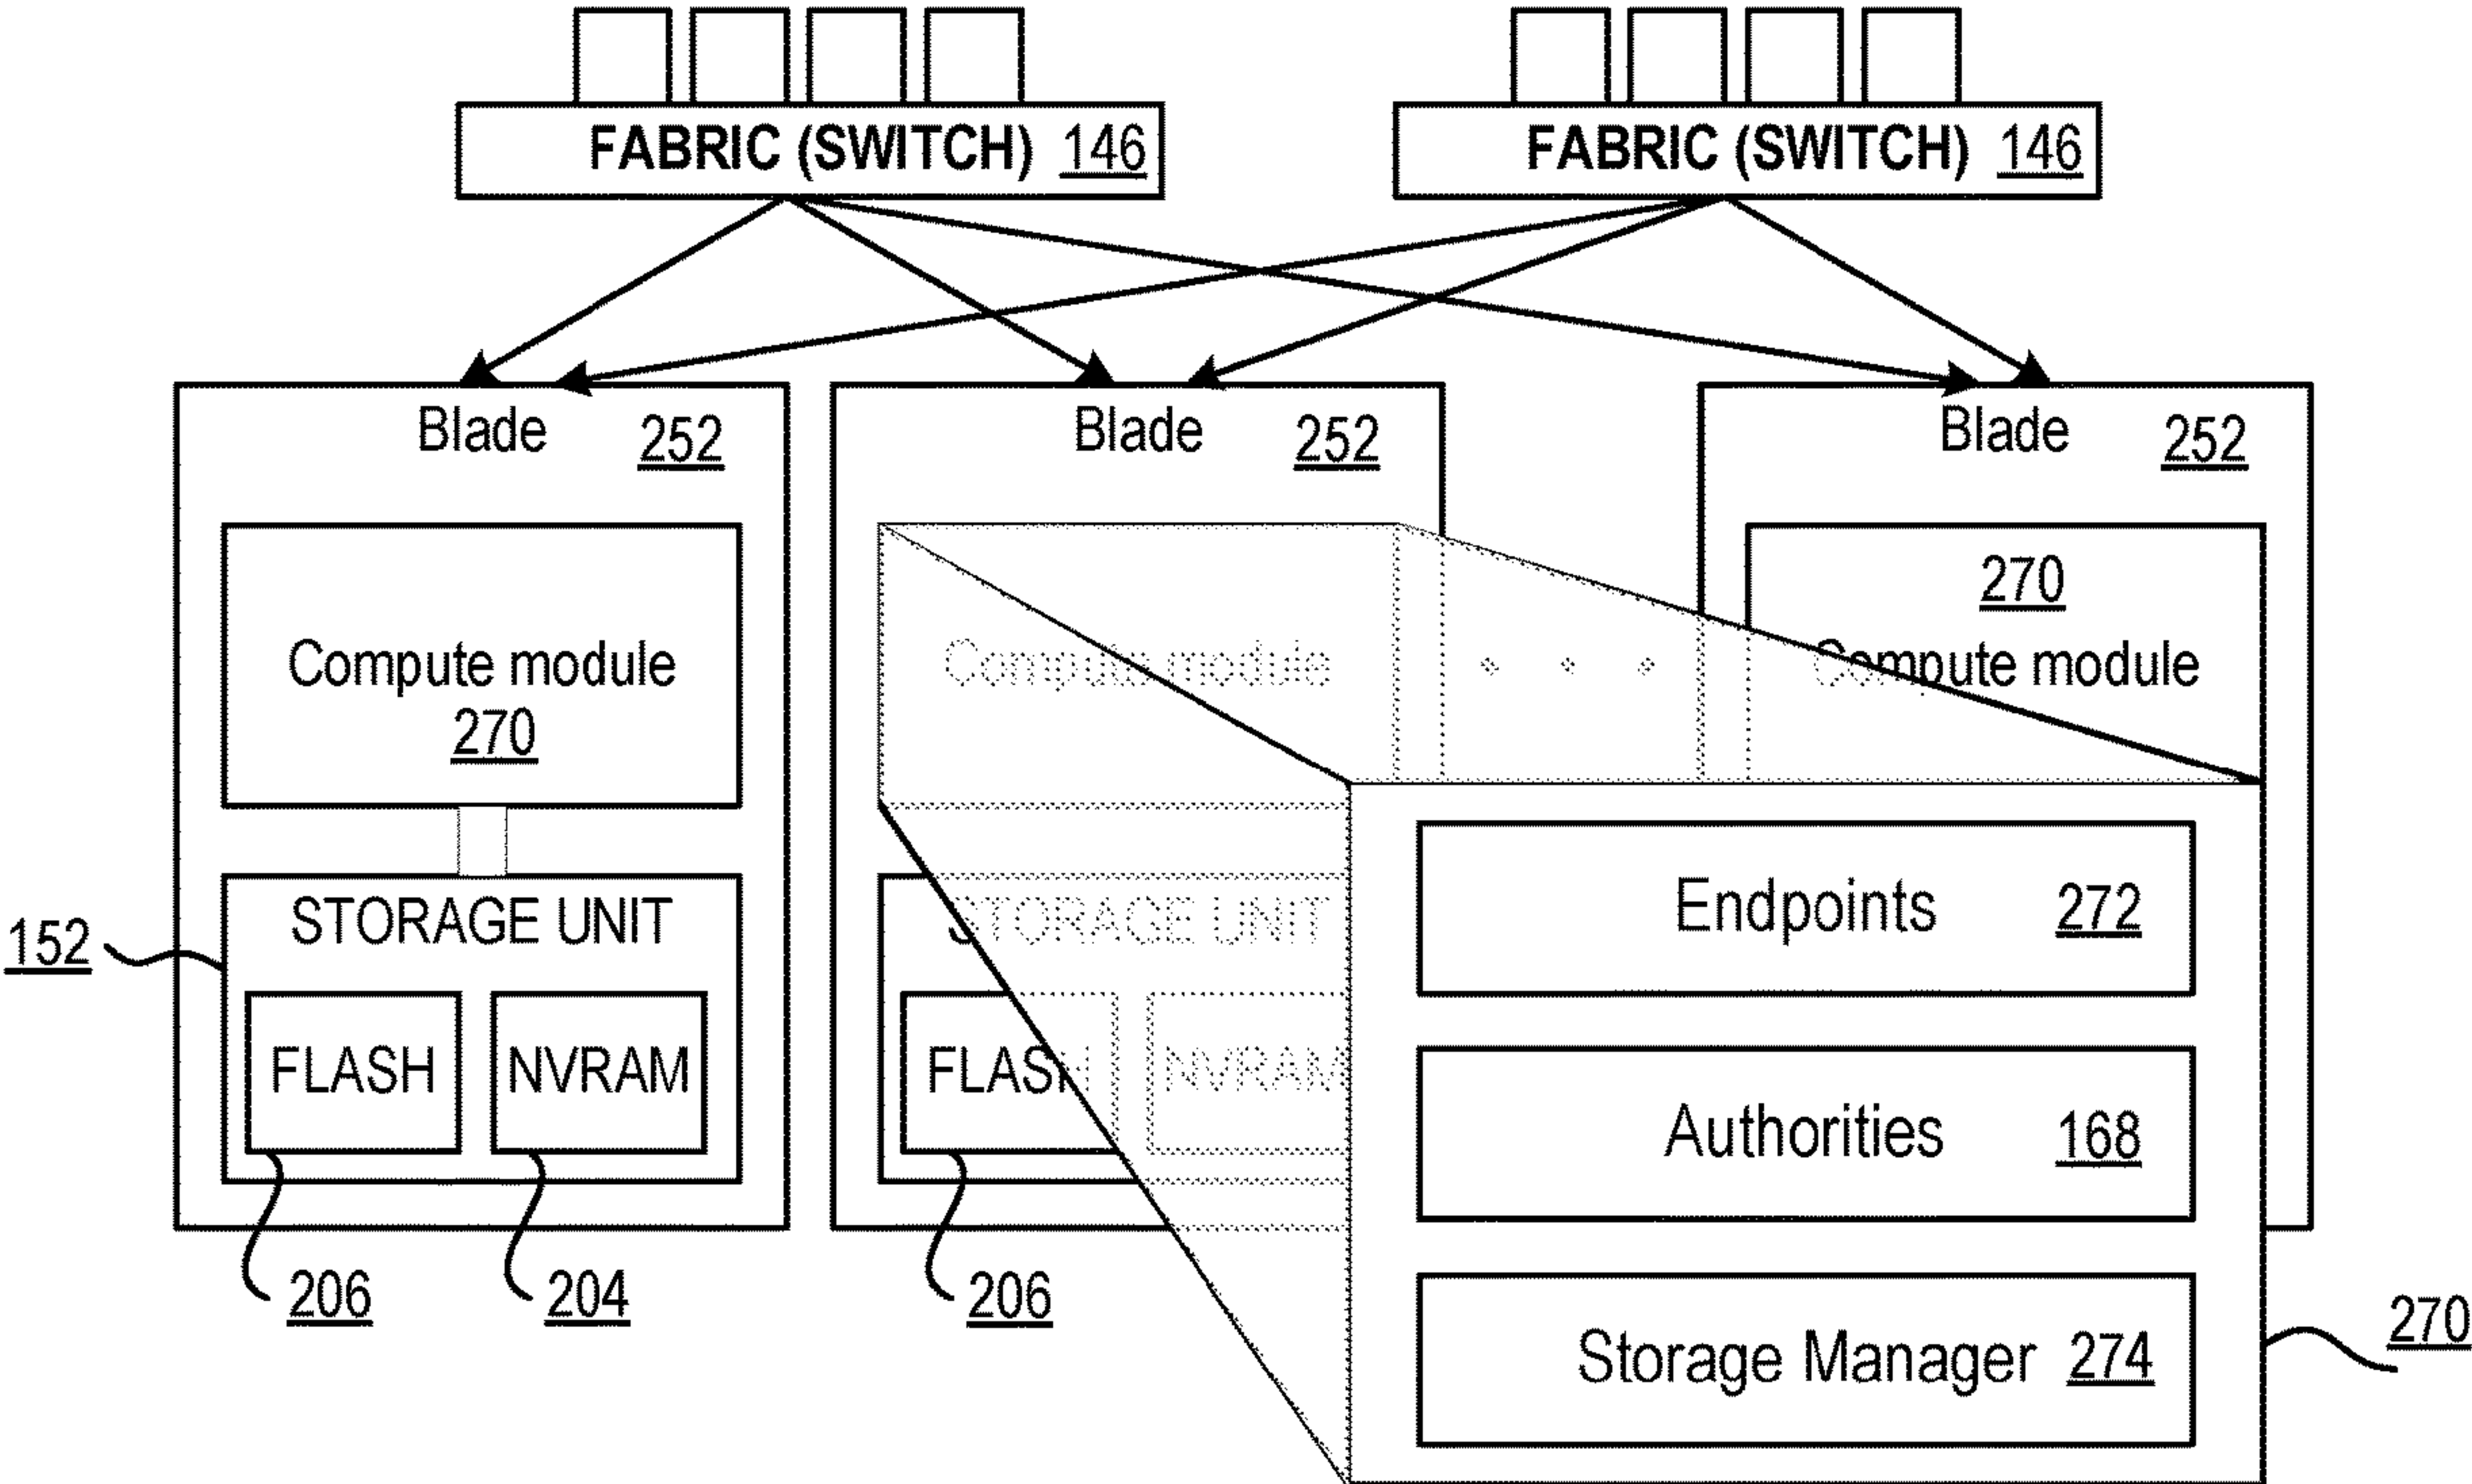


FIG. 2F




FIG. 2G




FIG. 3A




FIG. 3B




FIG. 3C

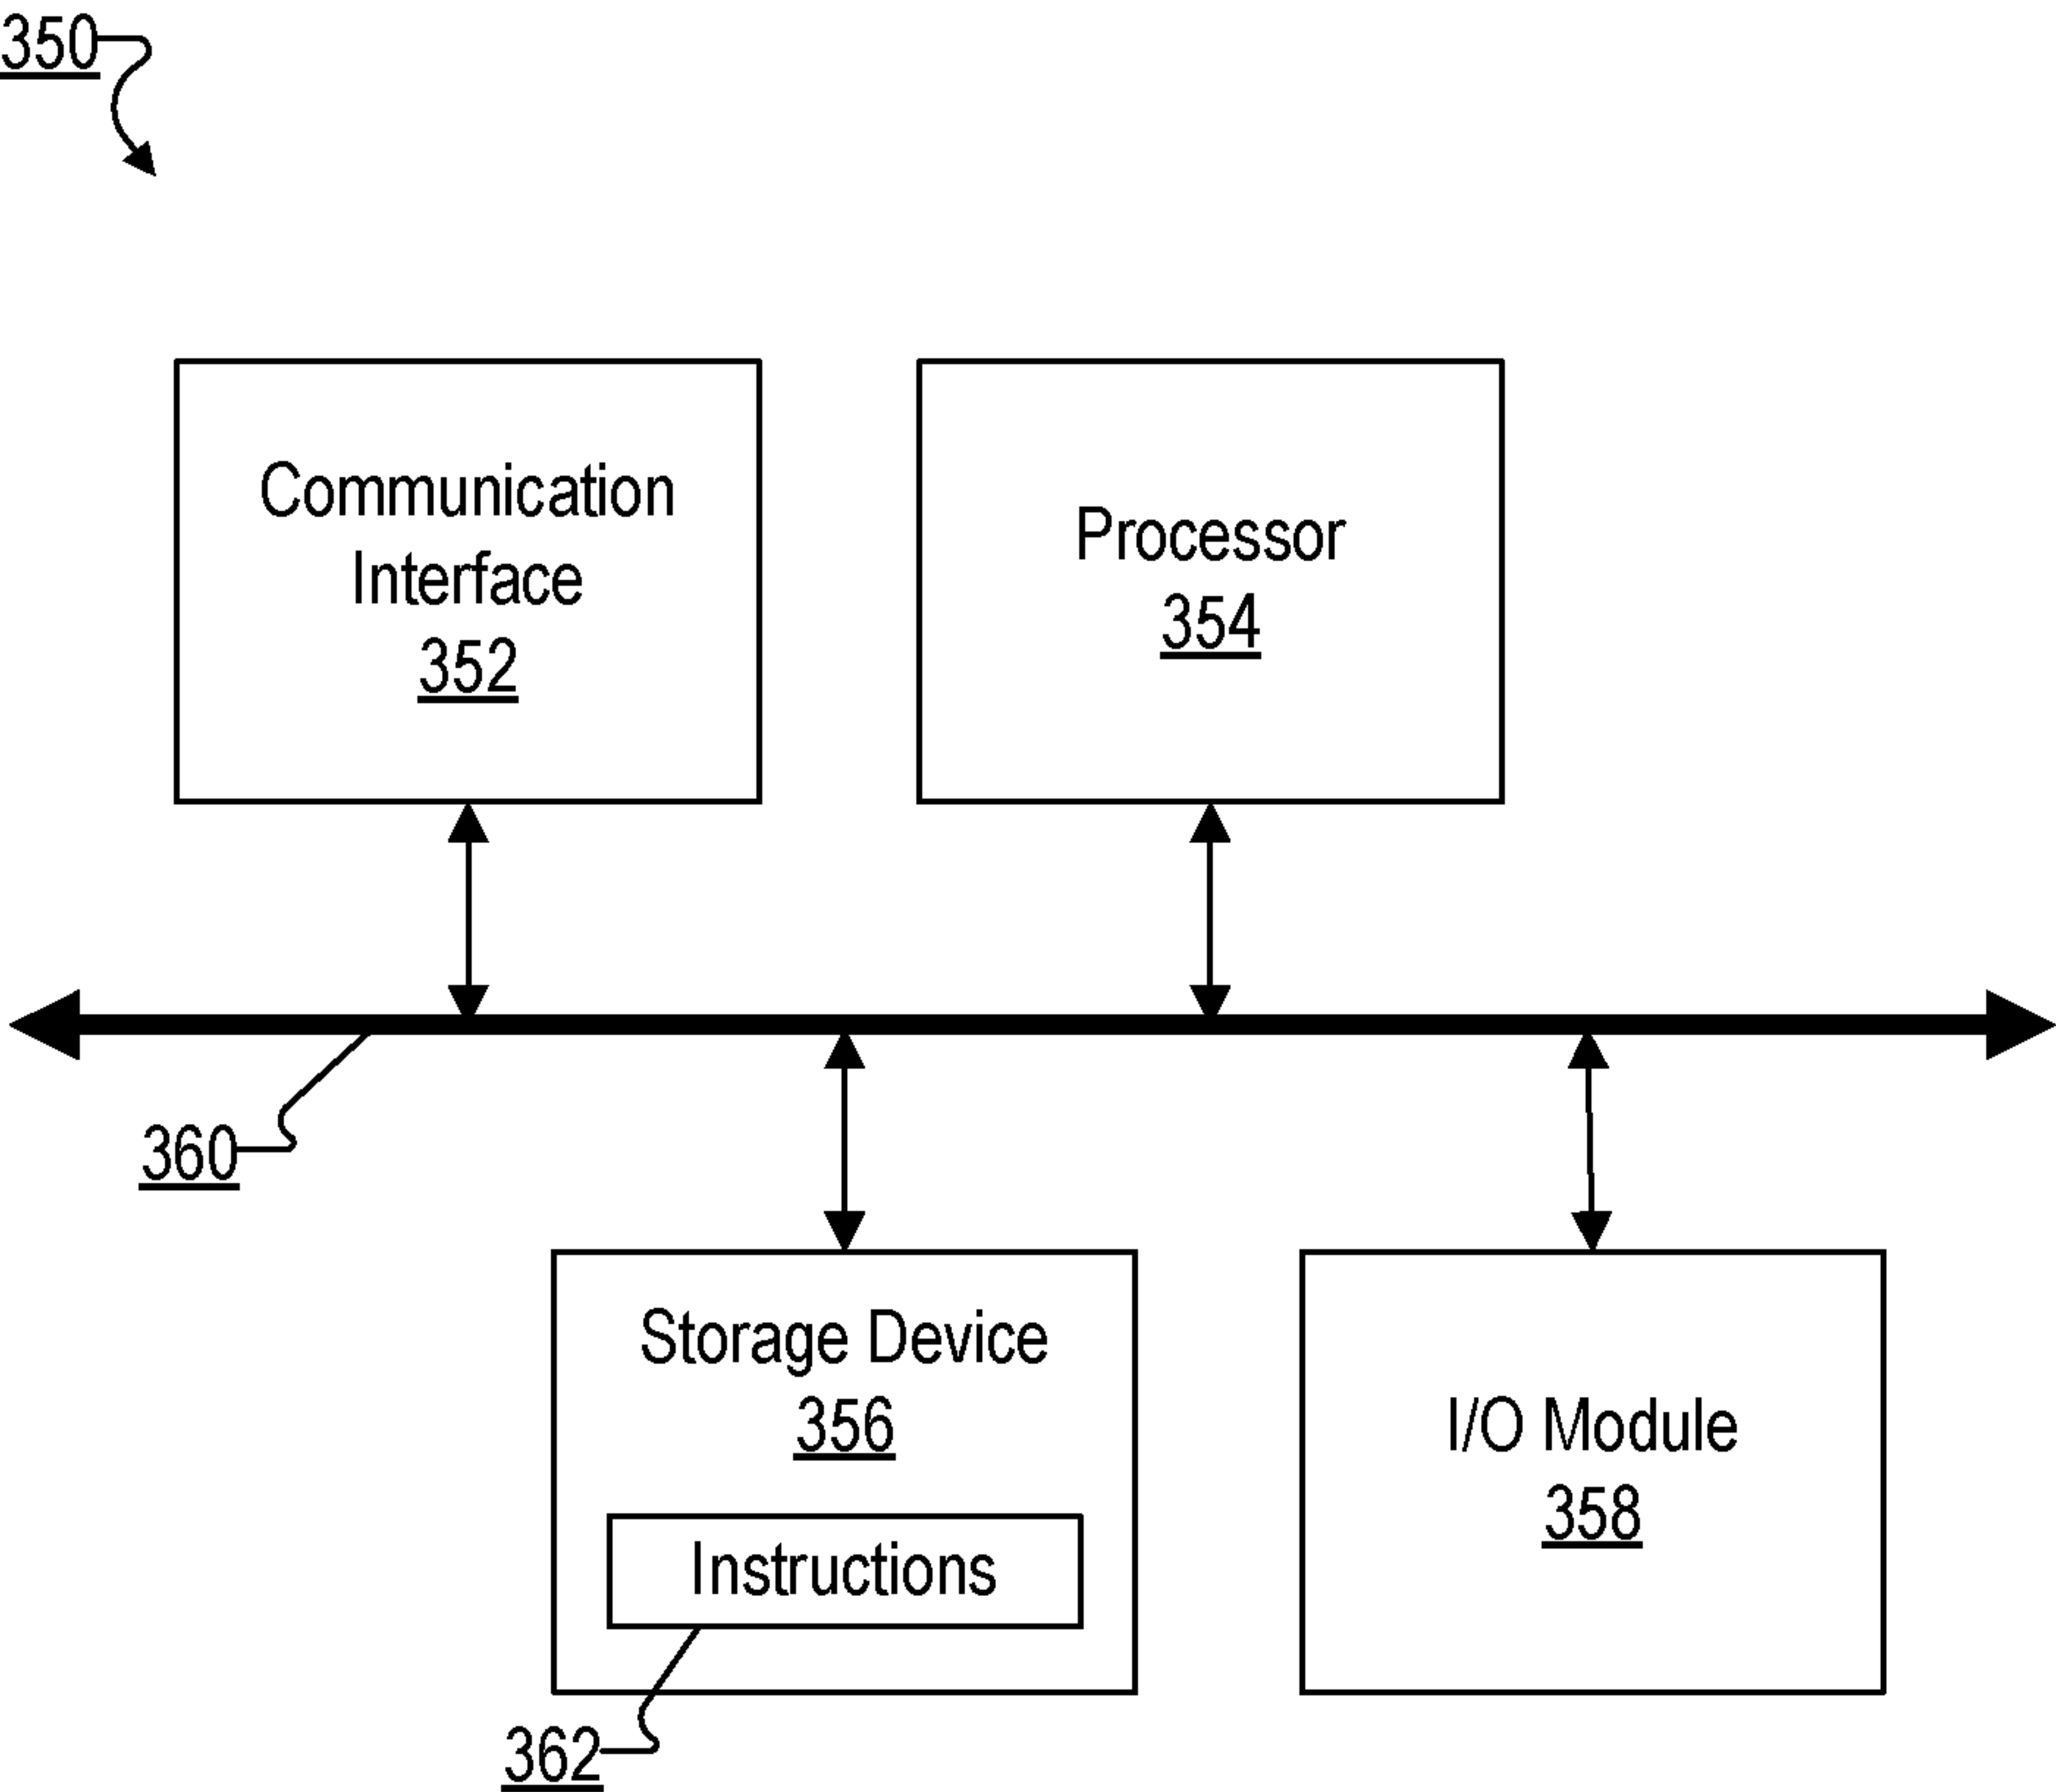


FIG. 3D

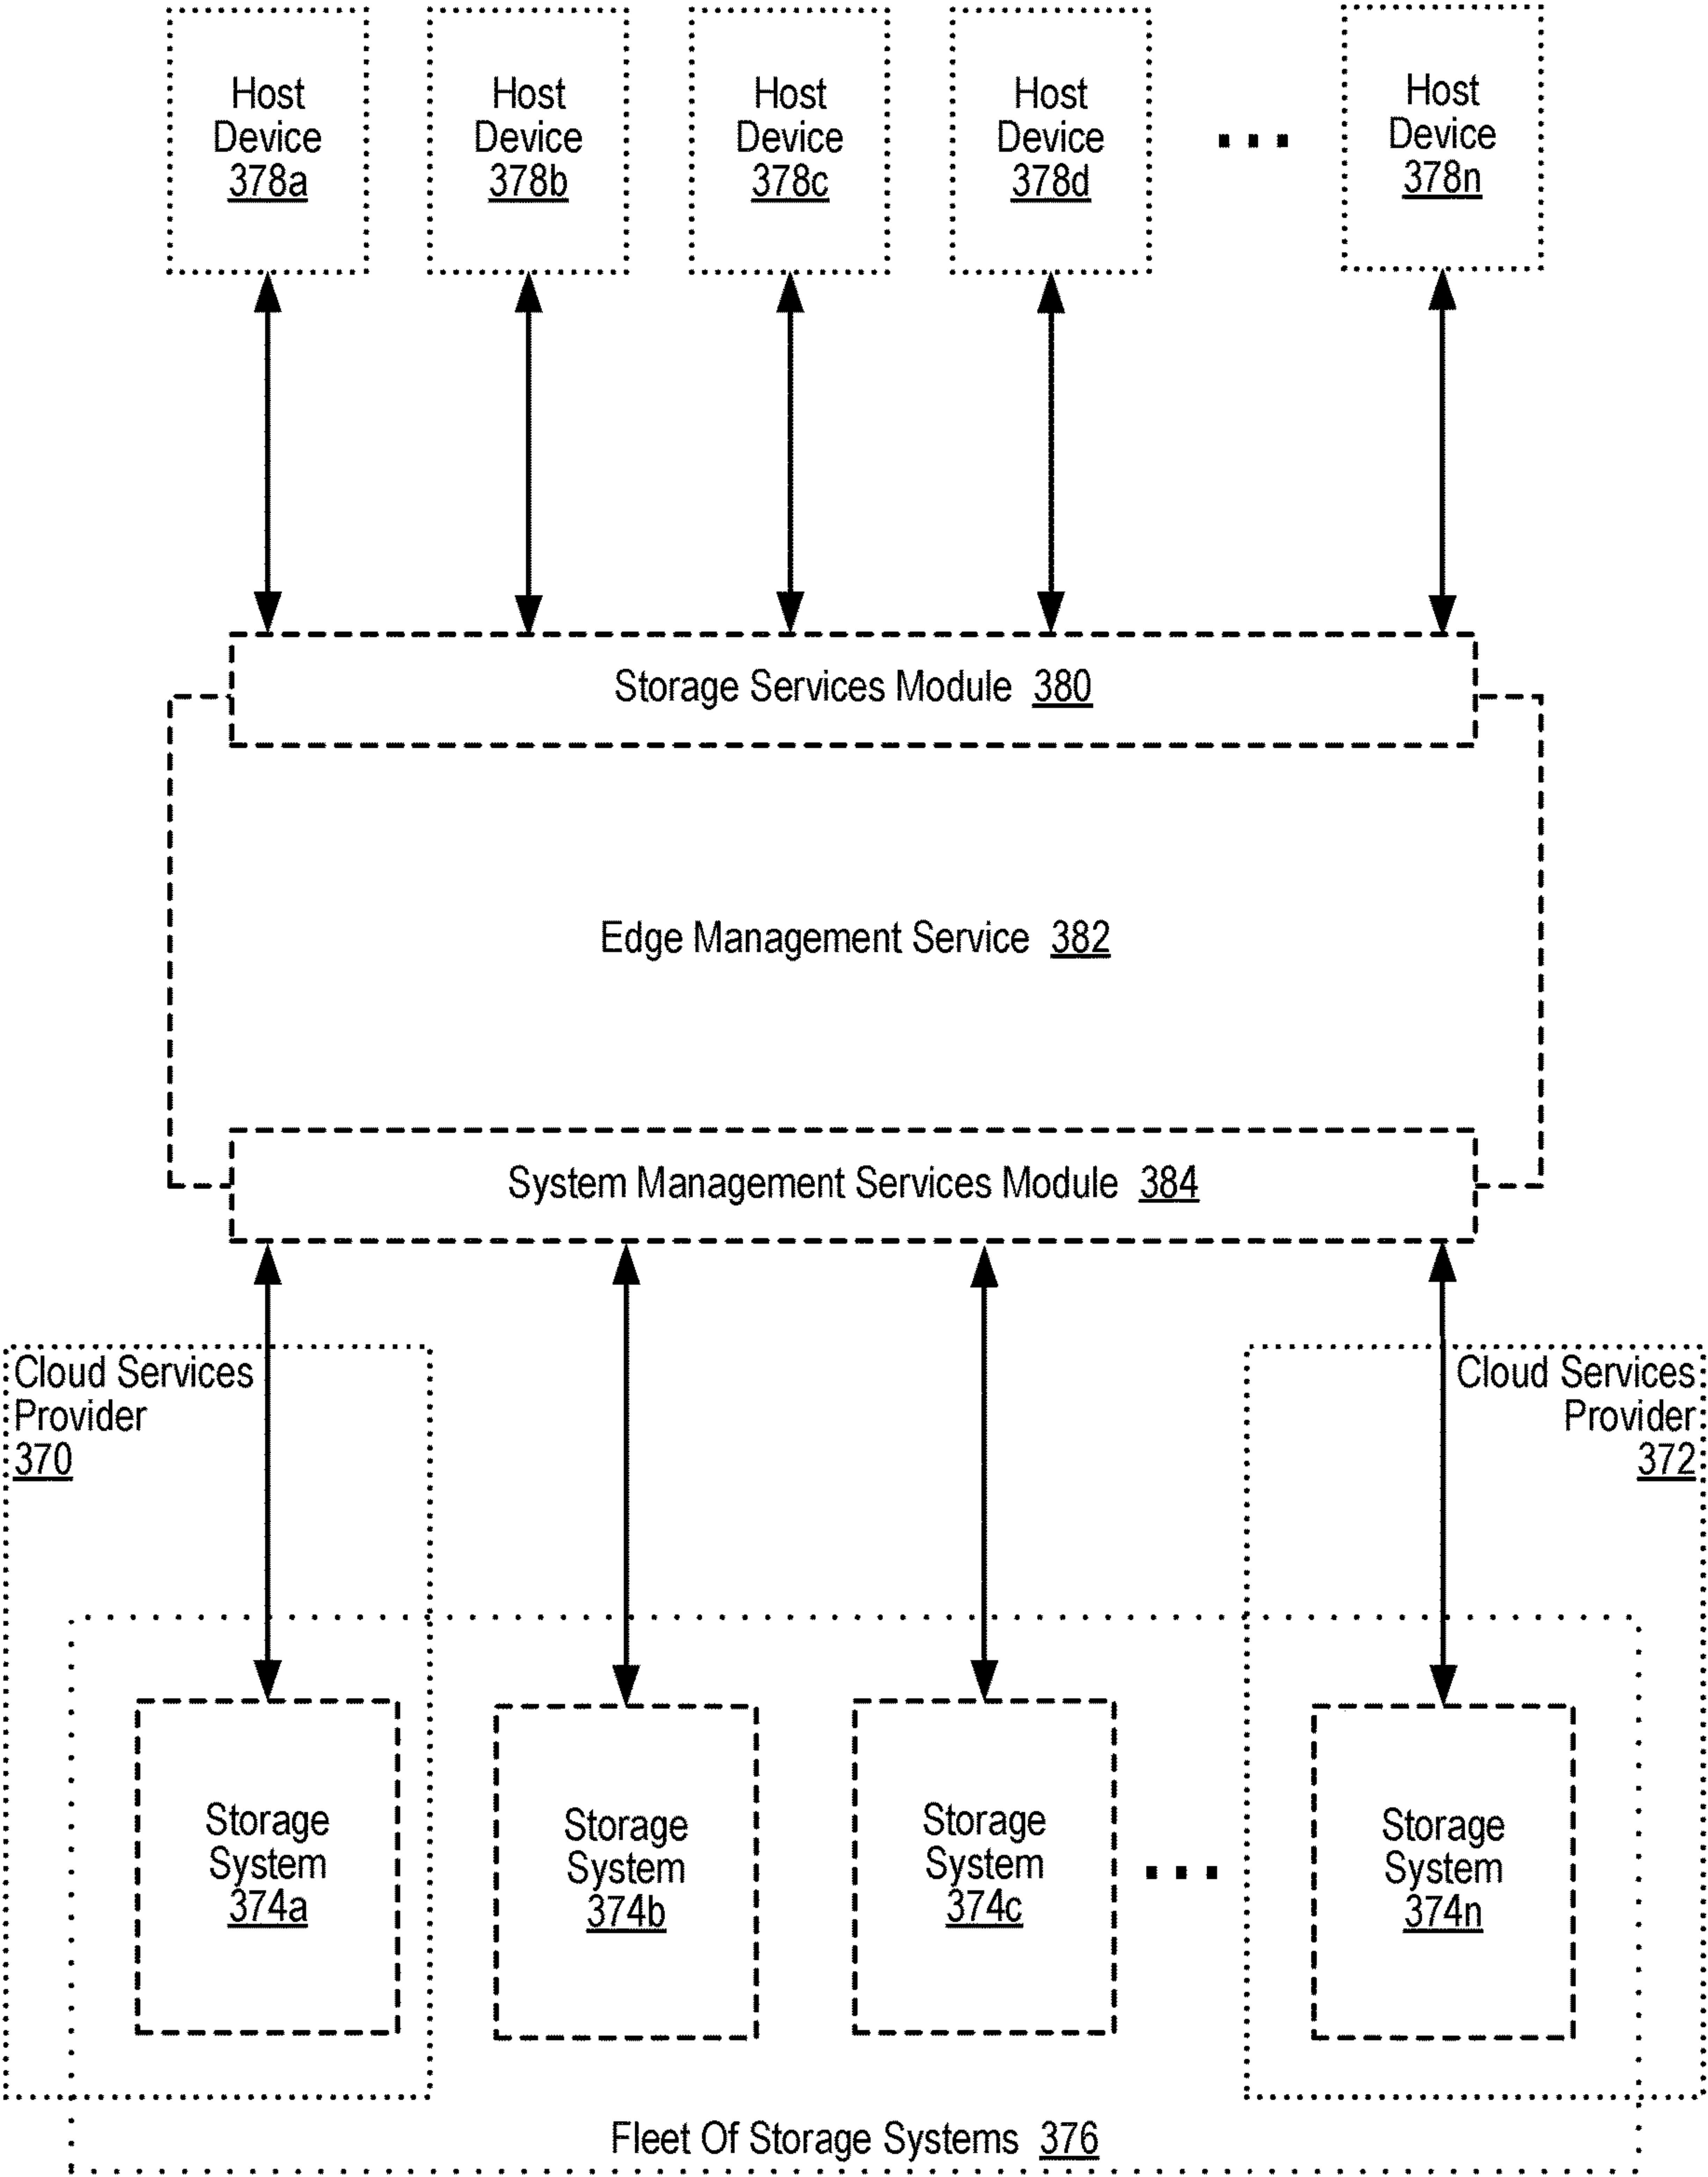


FIG. 3E

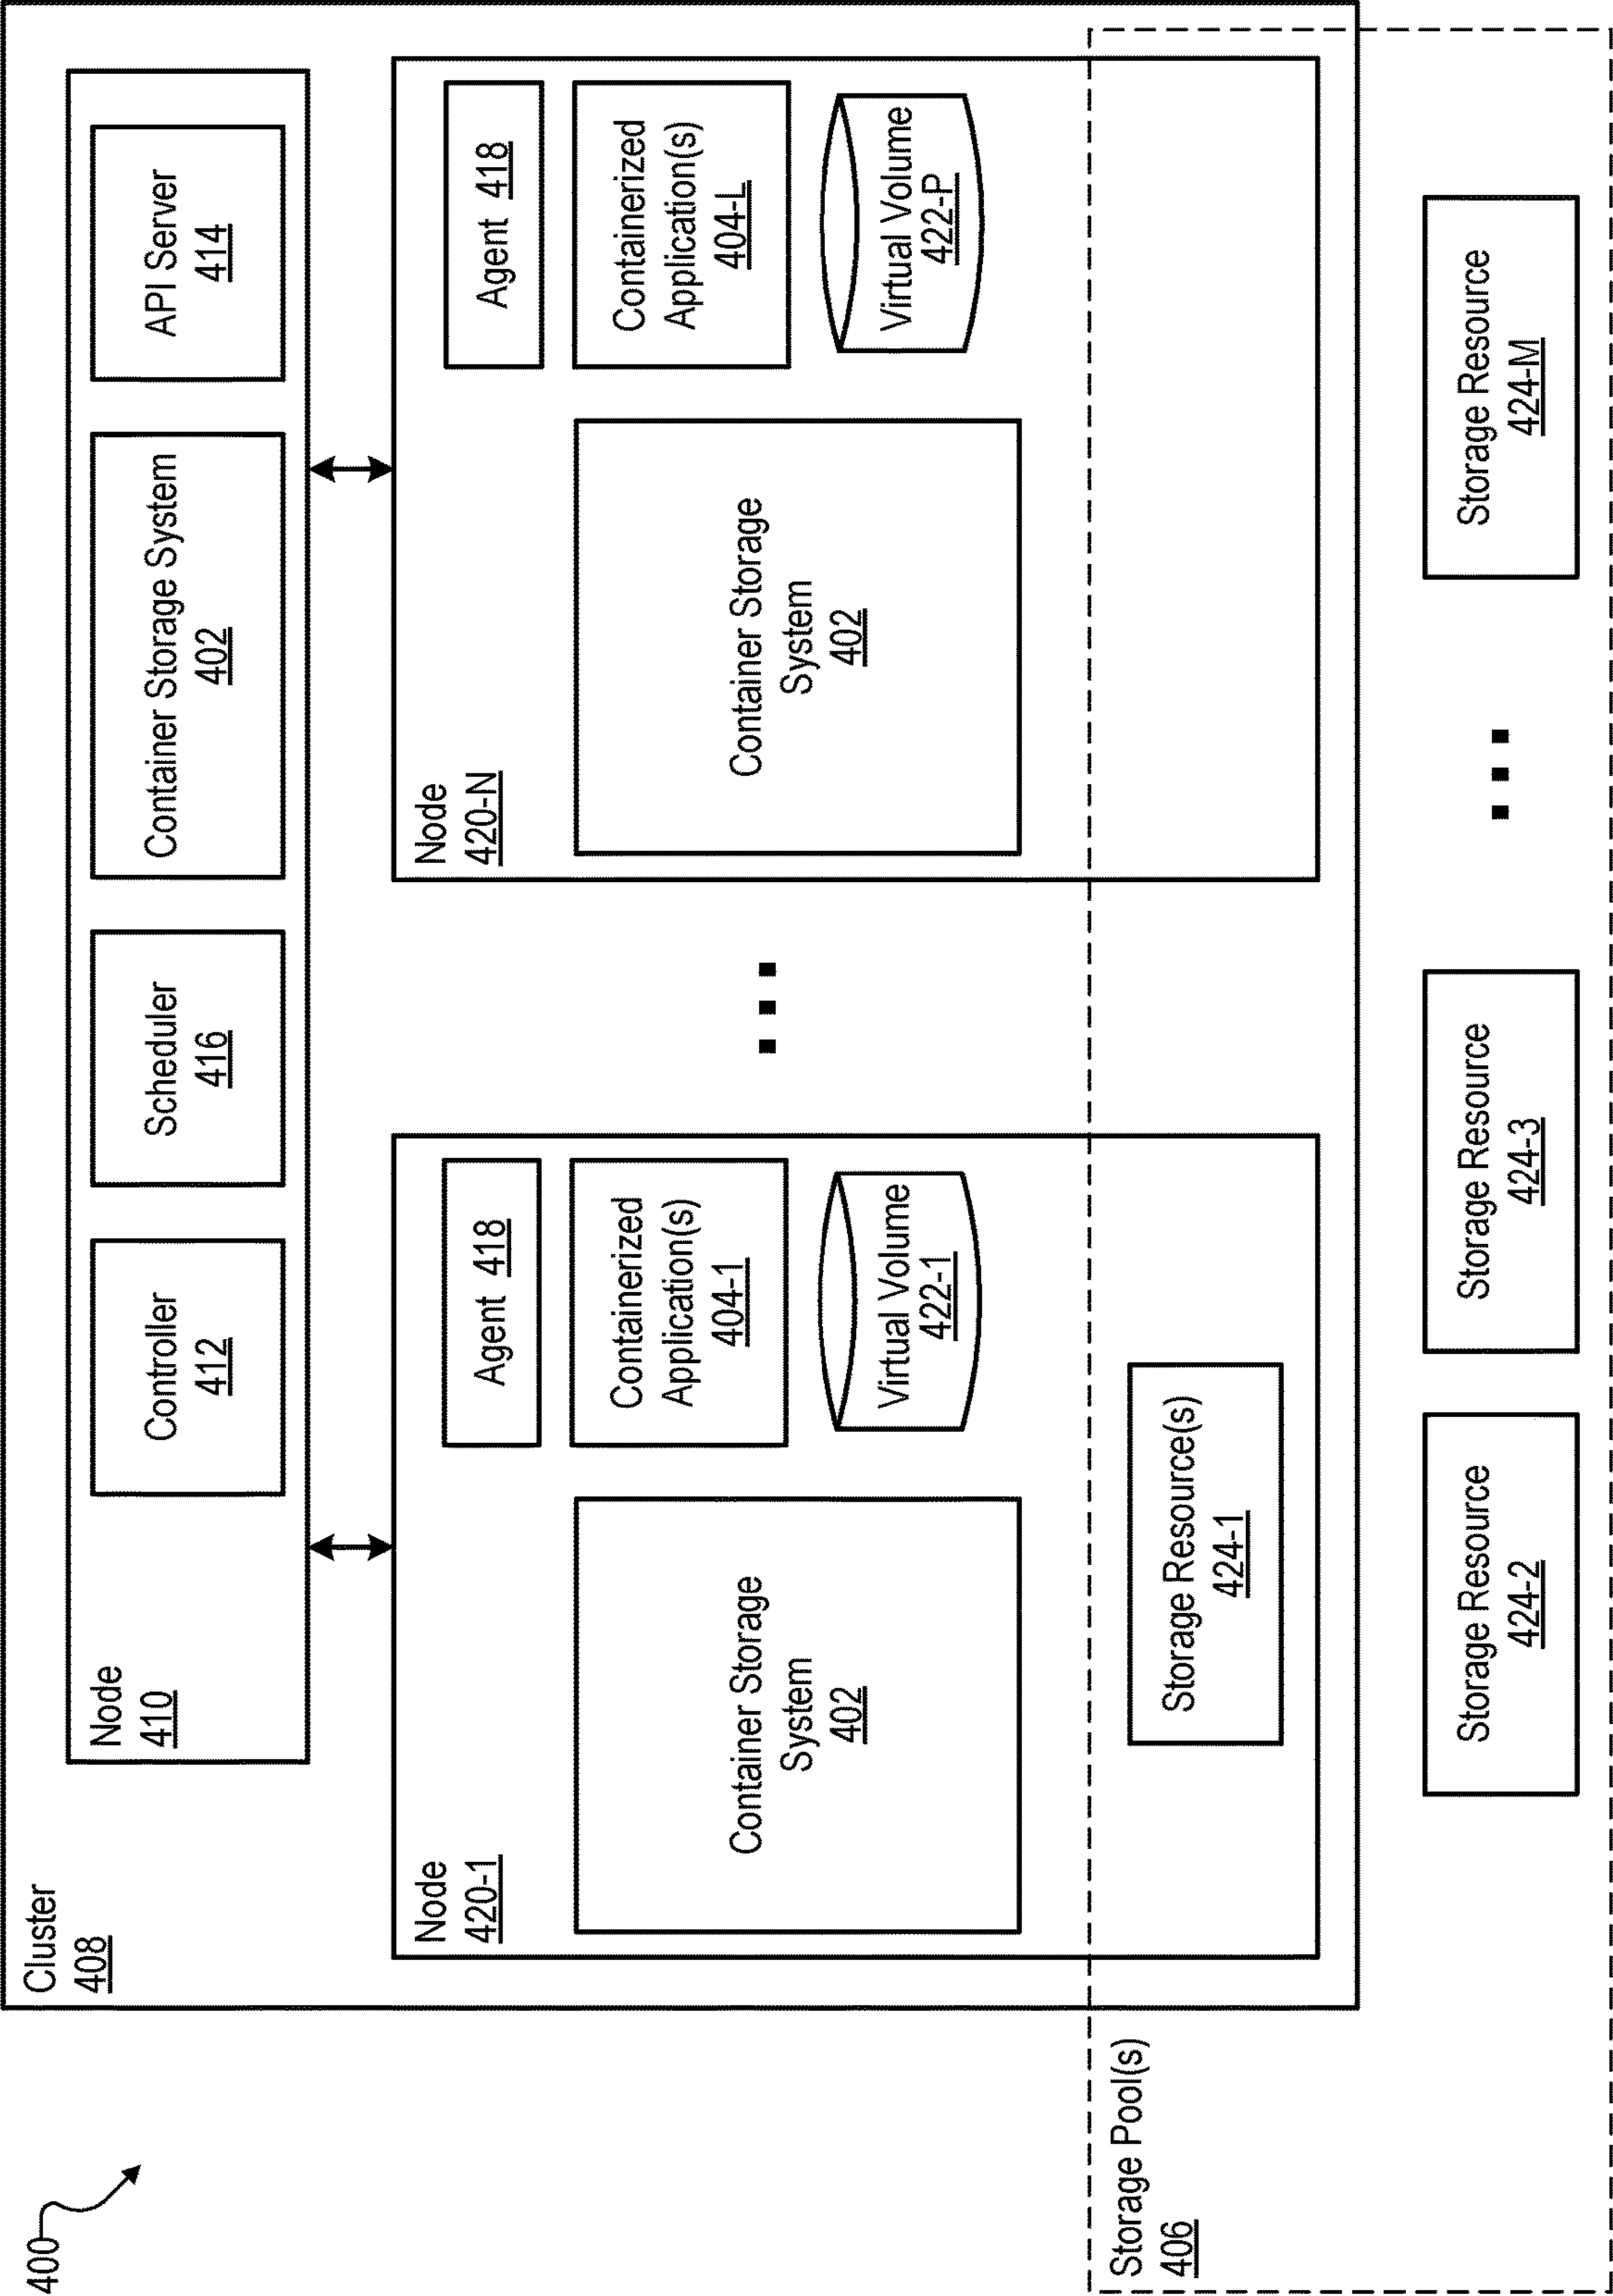
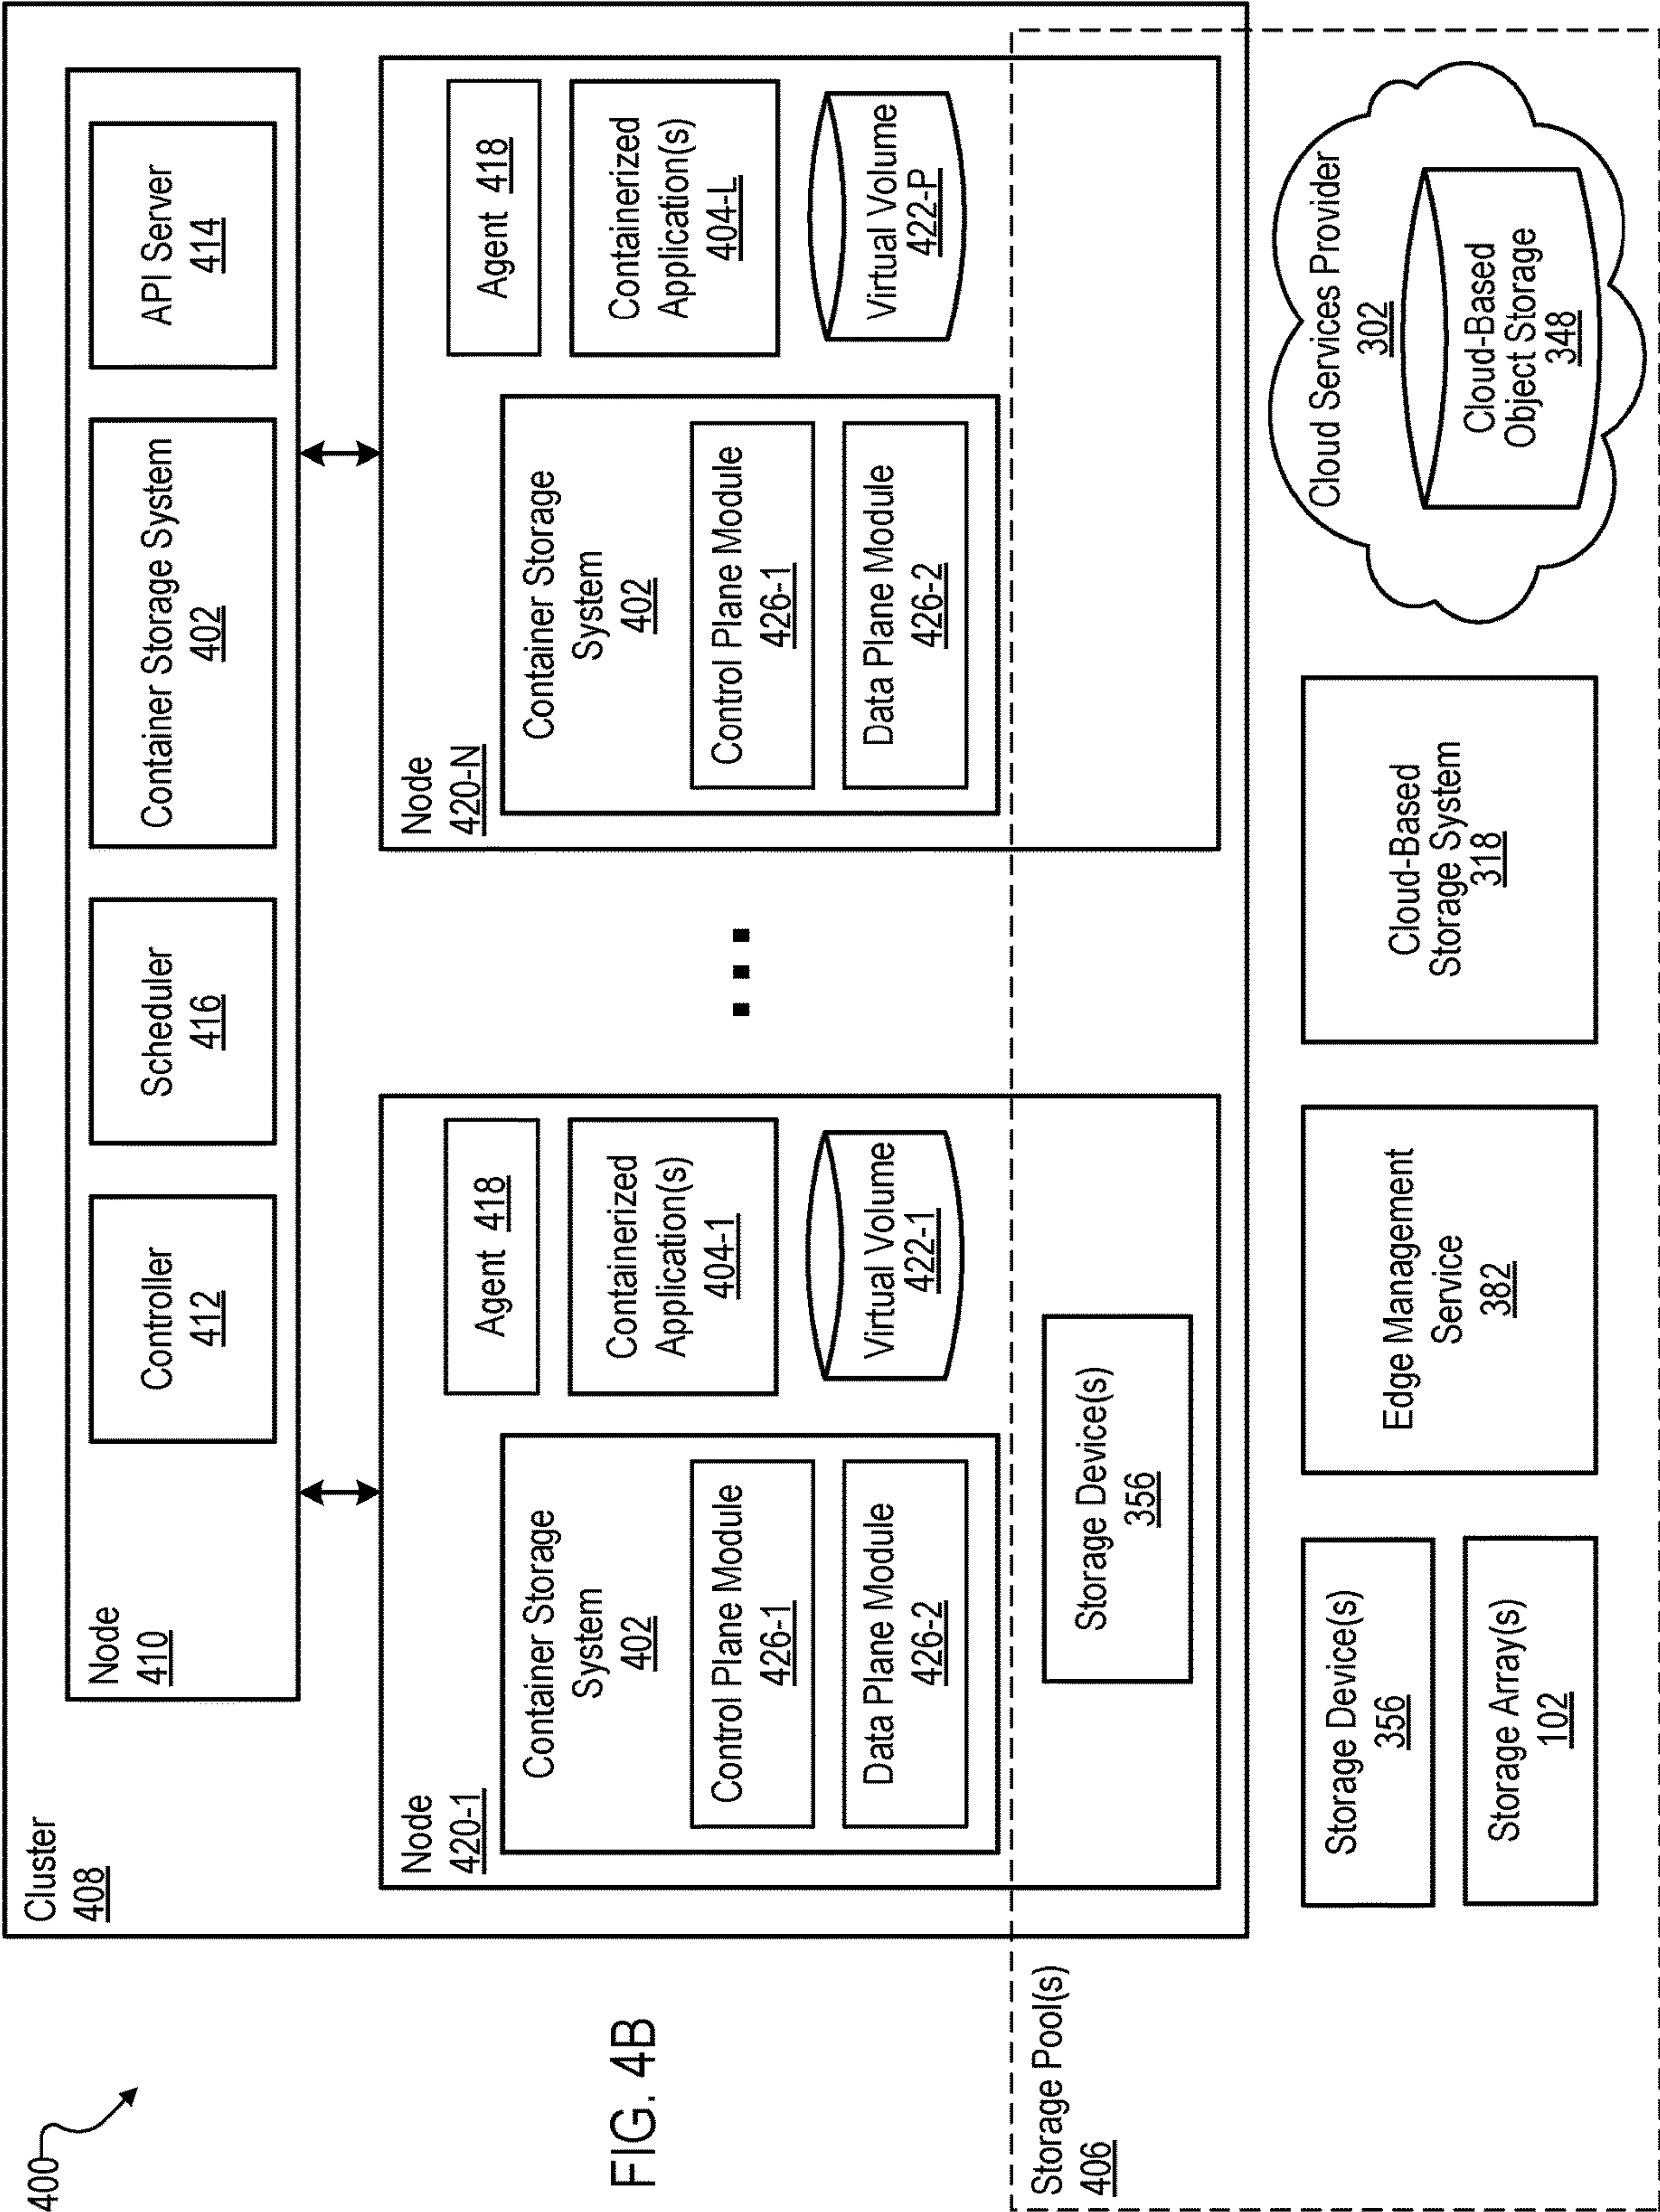


FIG. 4A



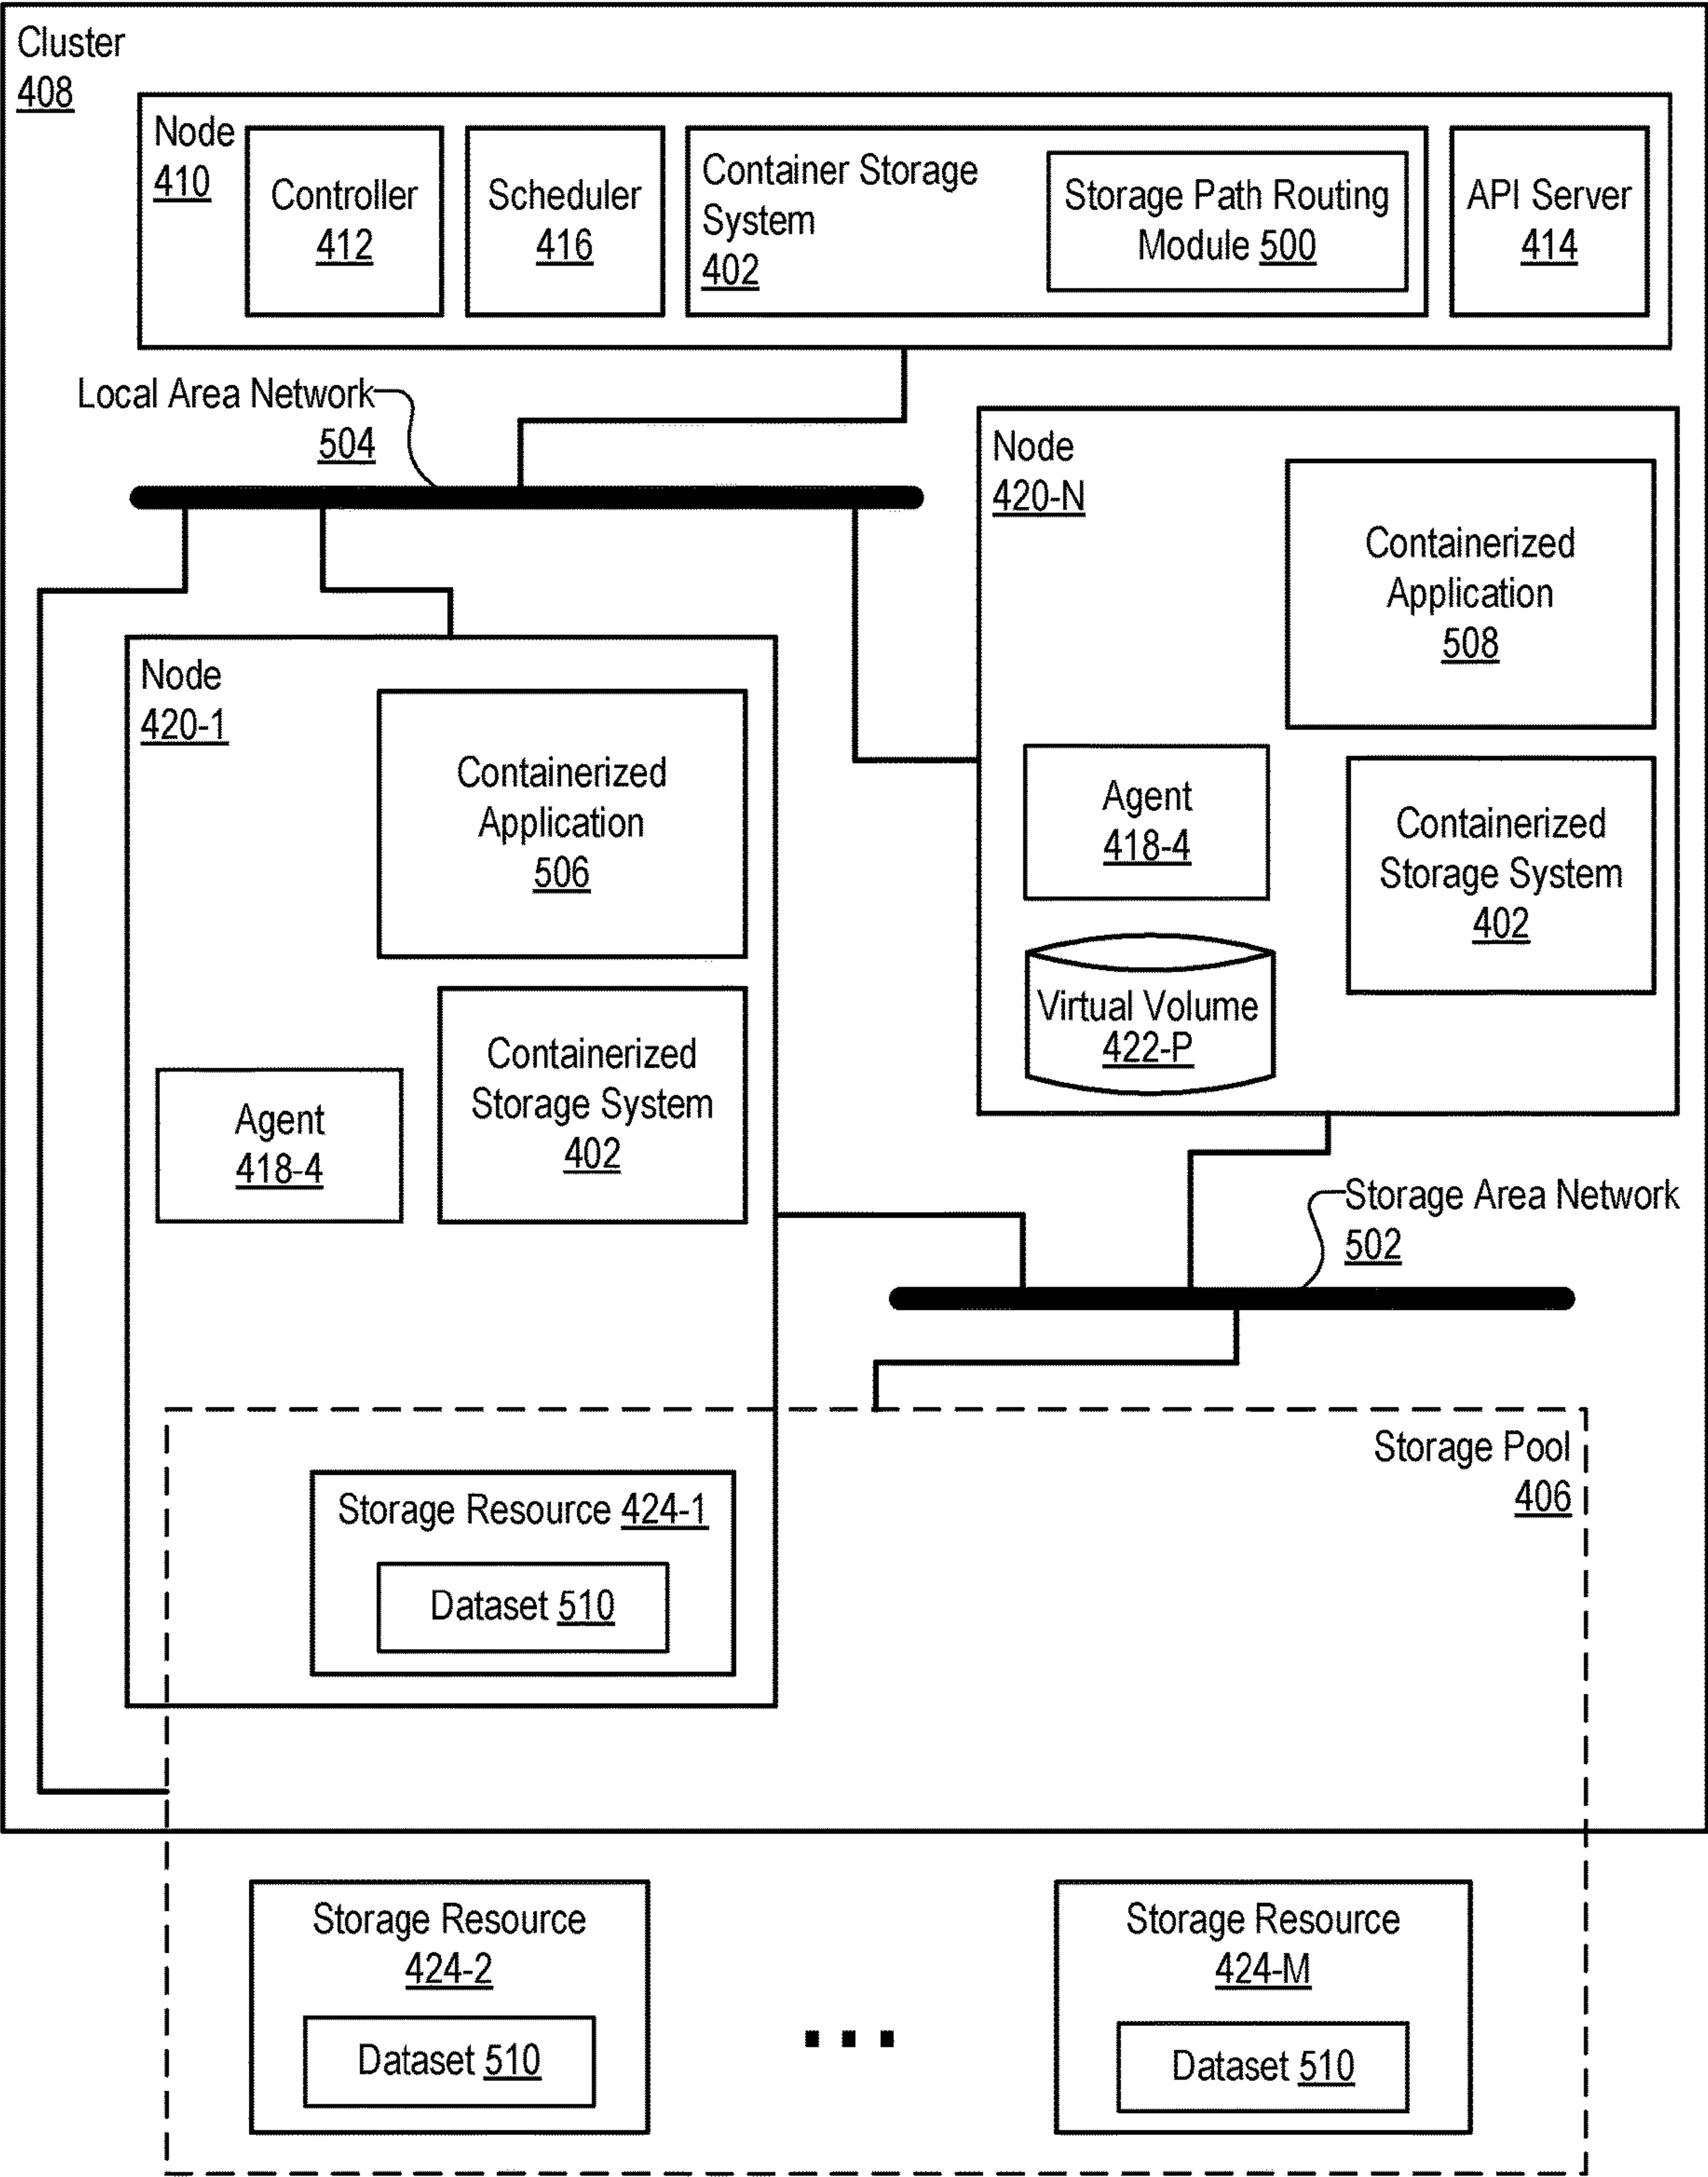


FIG. 5

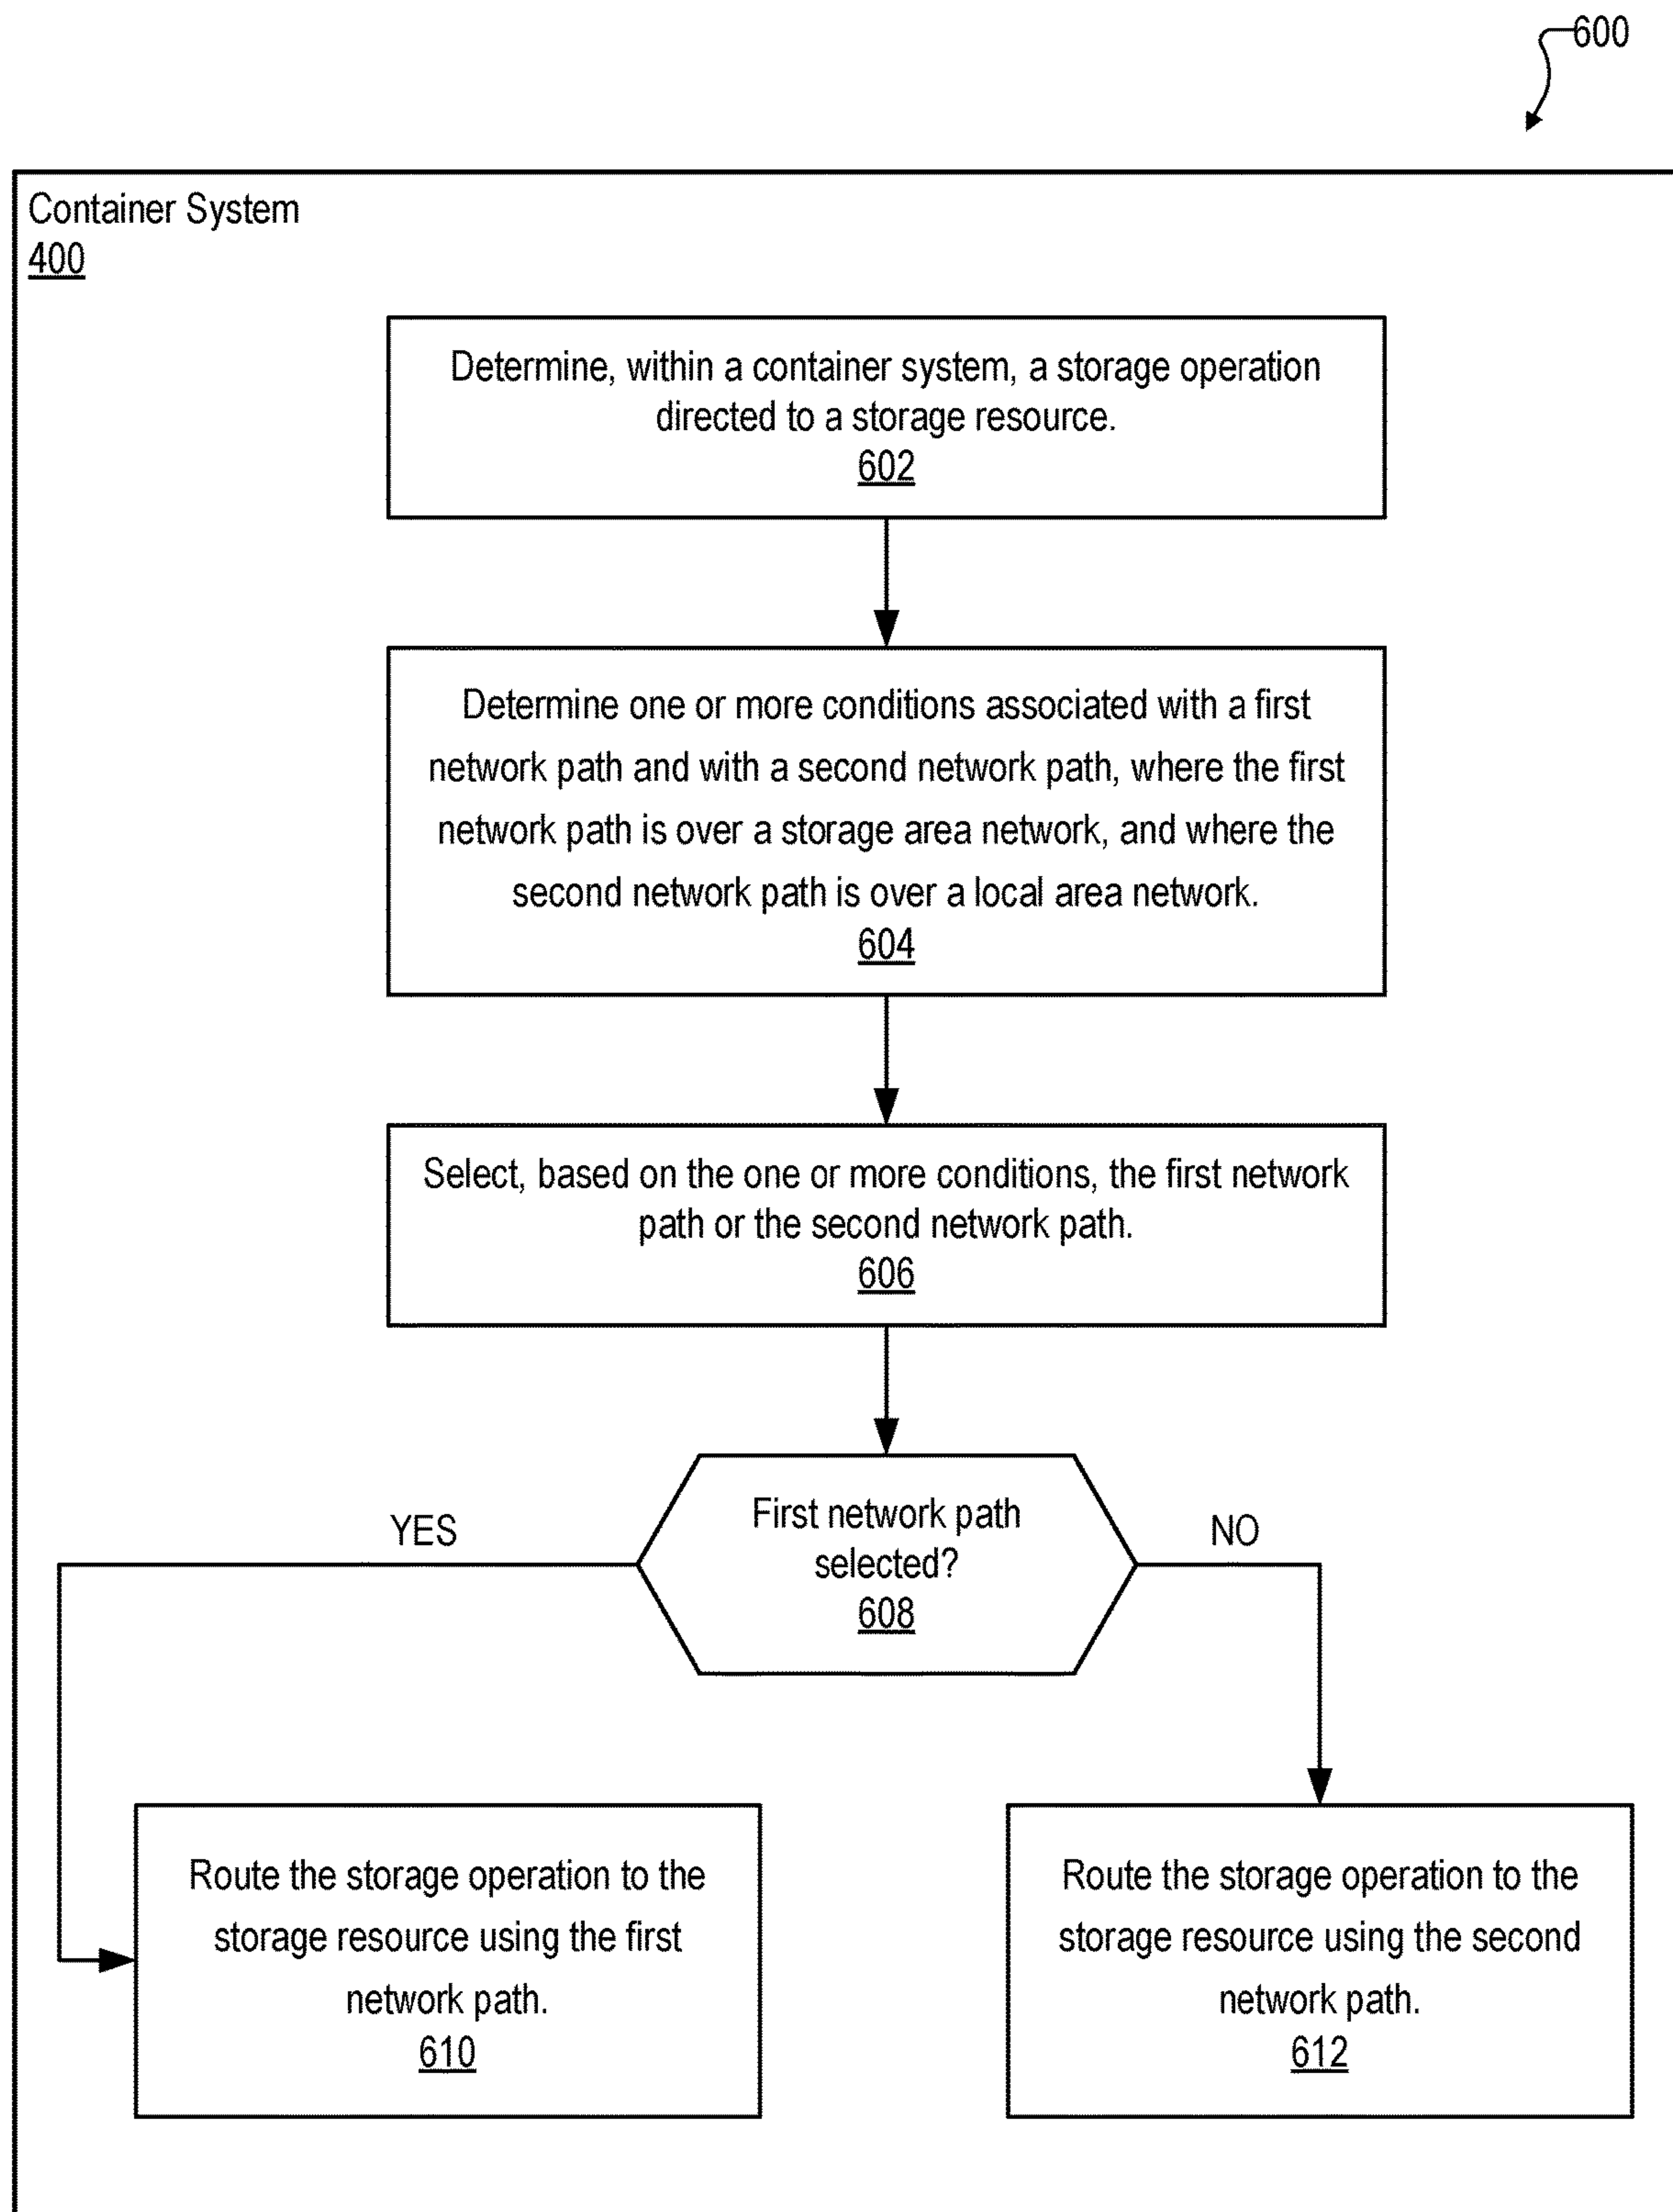


FIG. 6

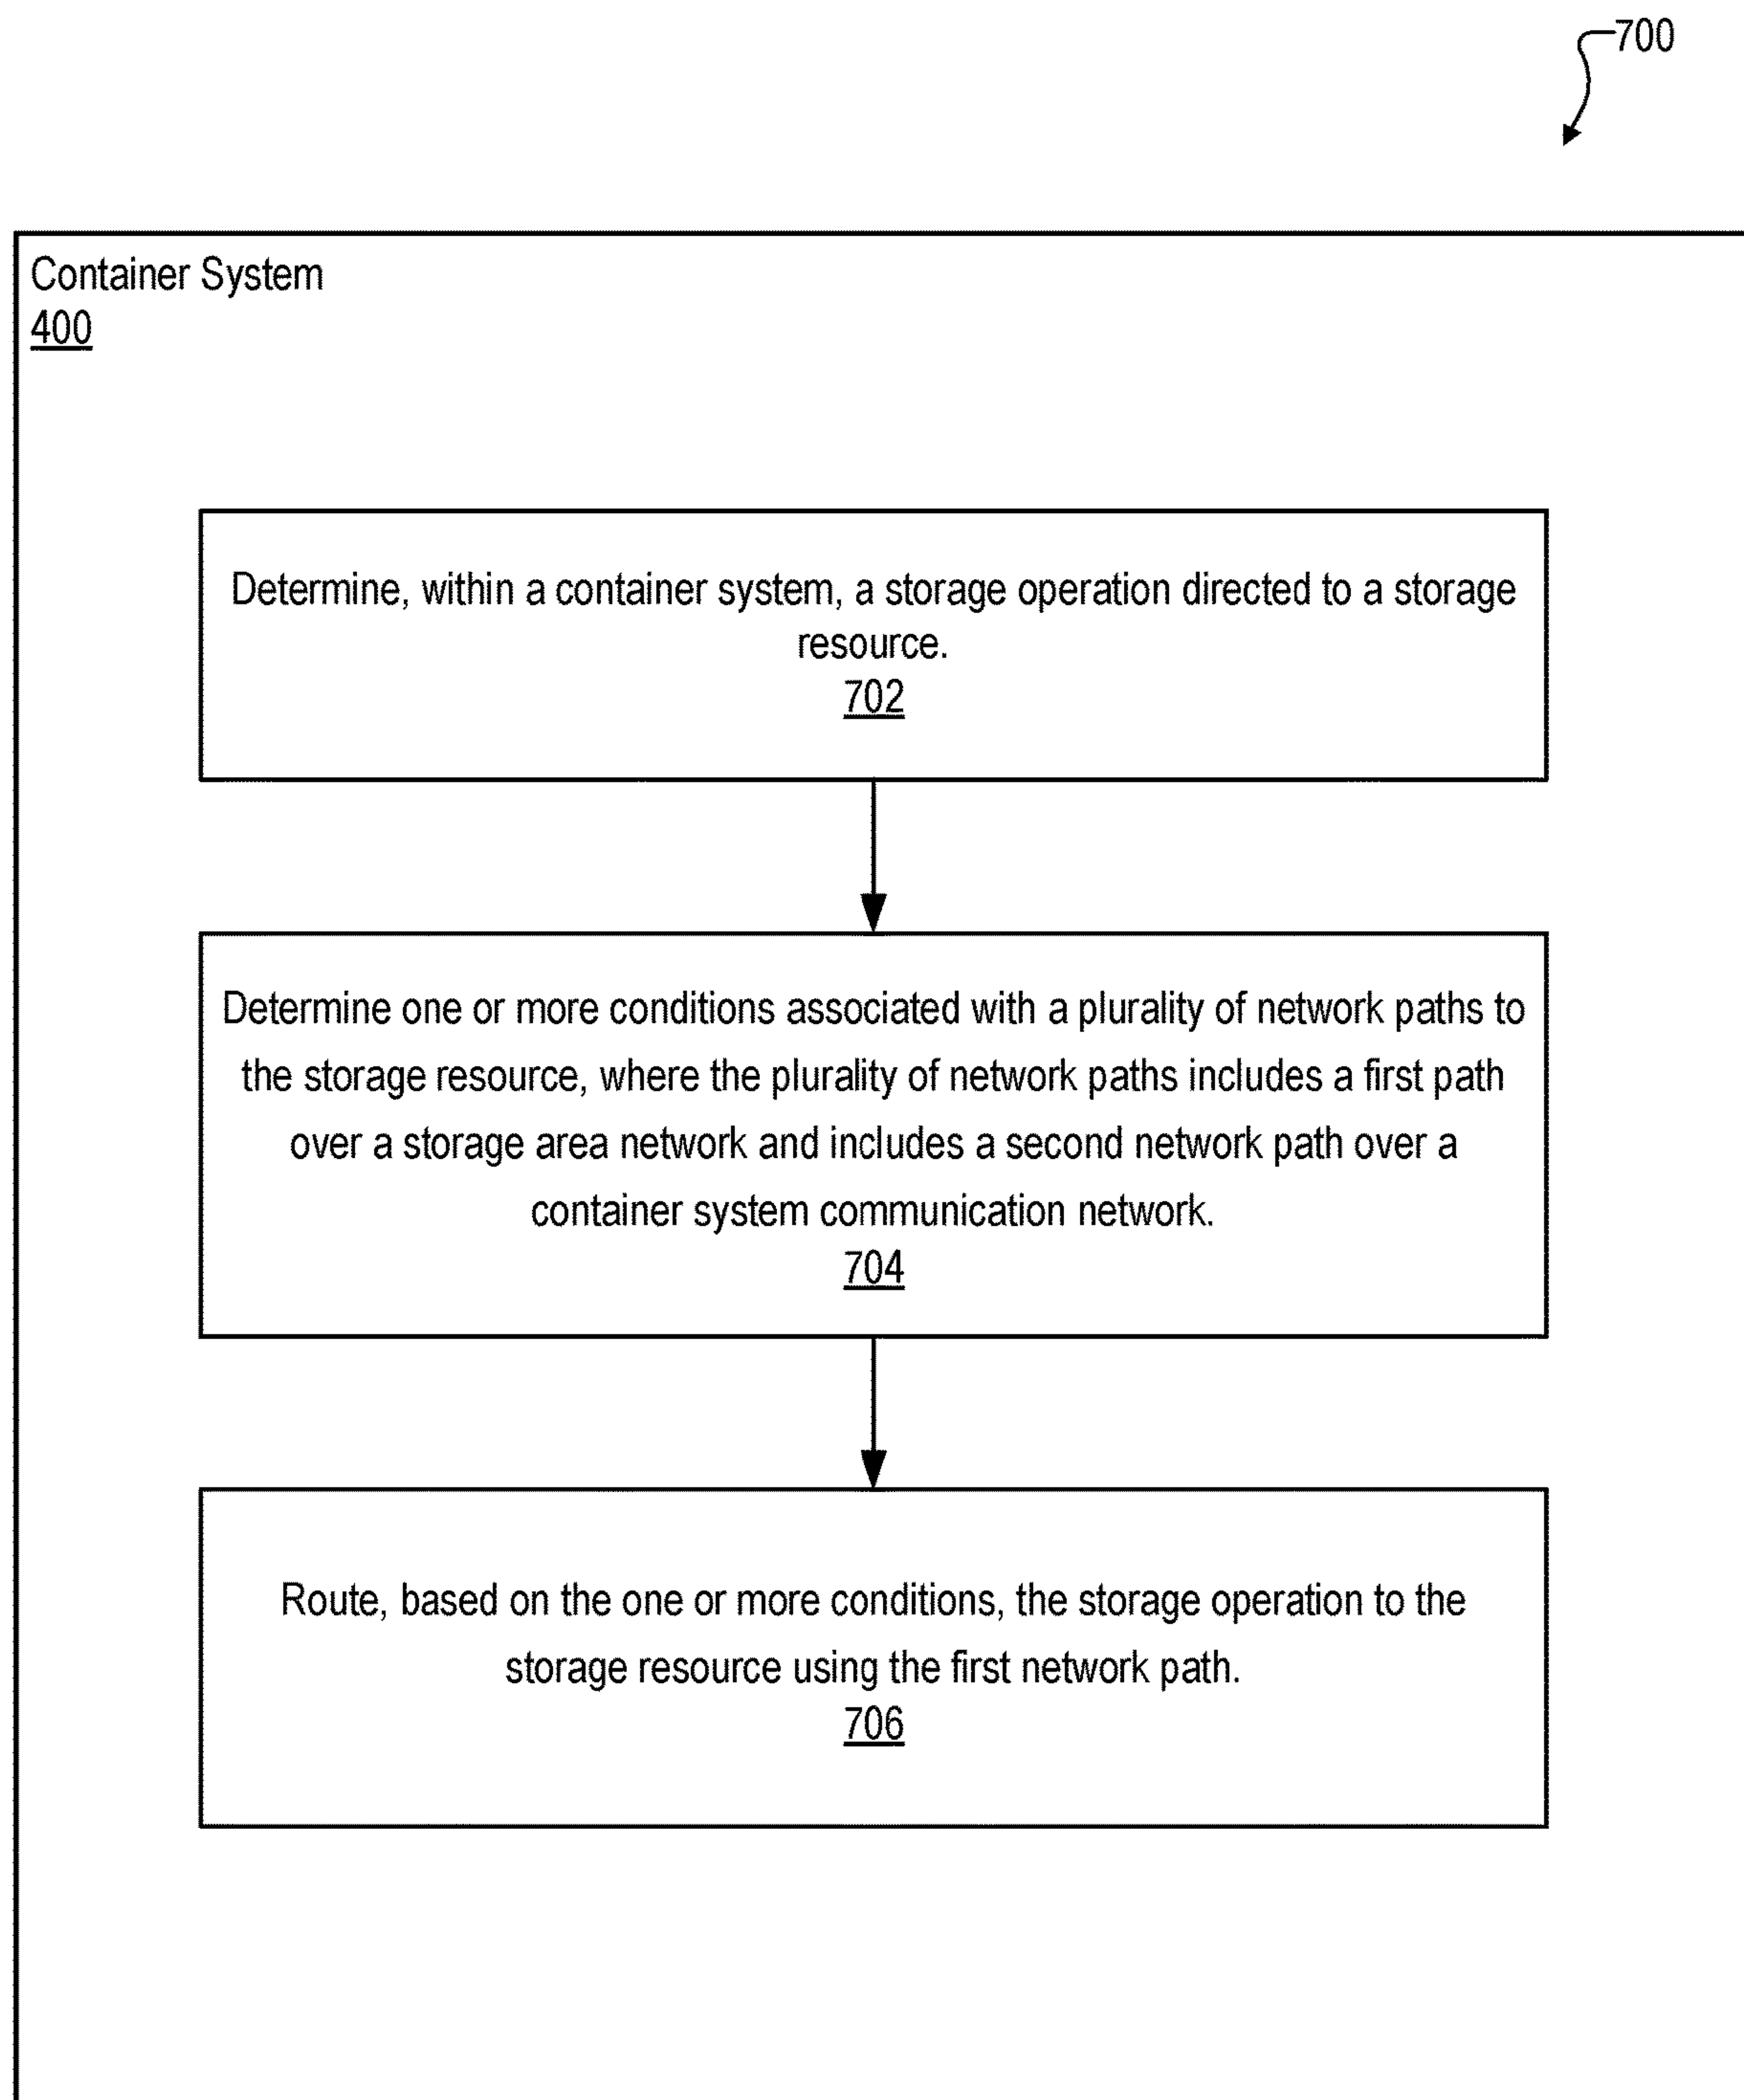


FIG. 7

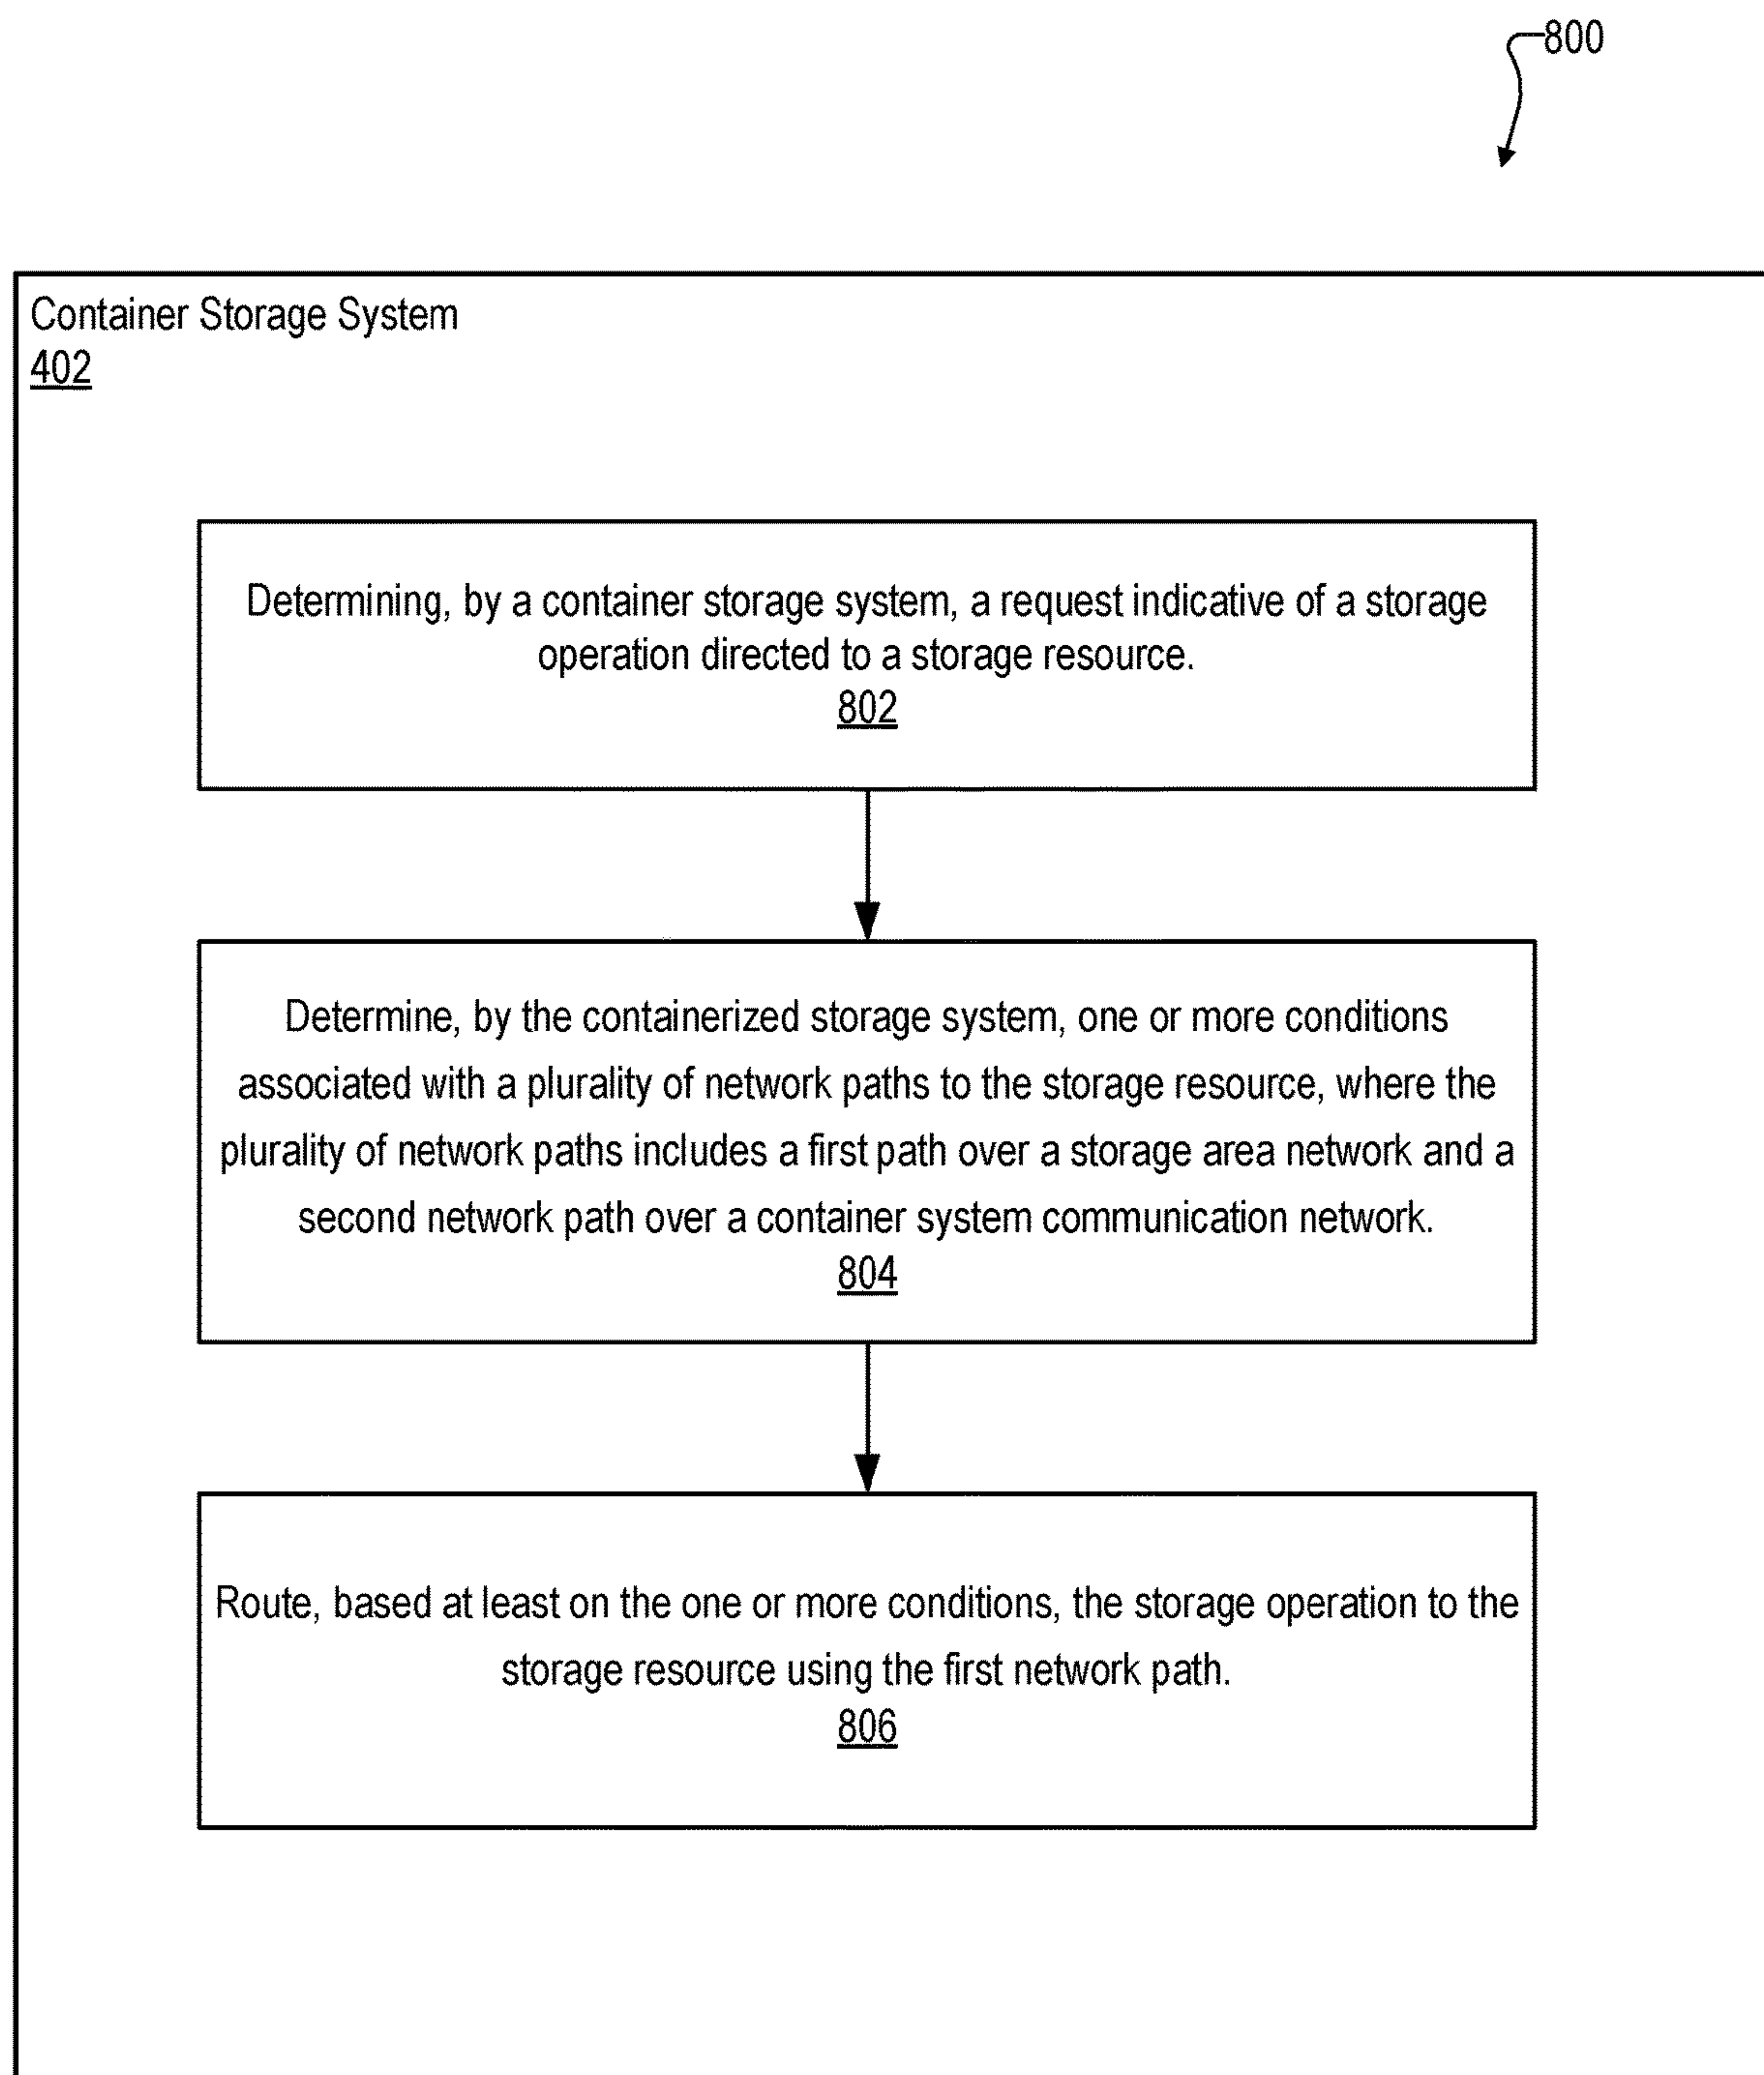


FIG. 8

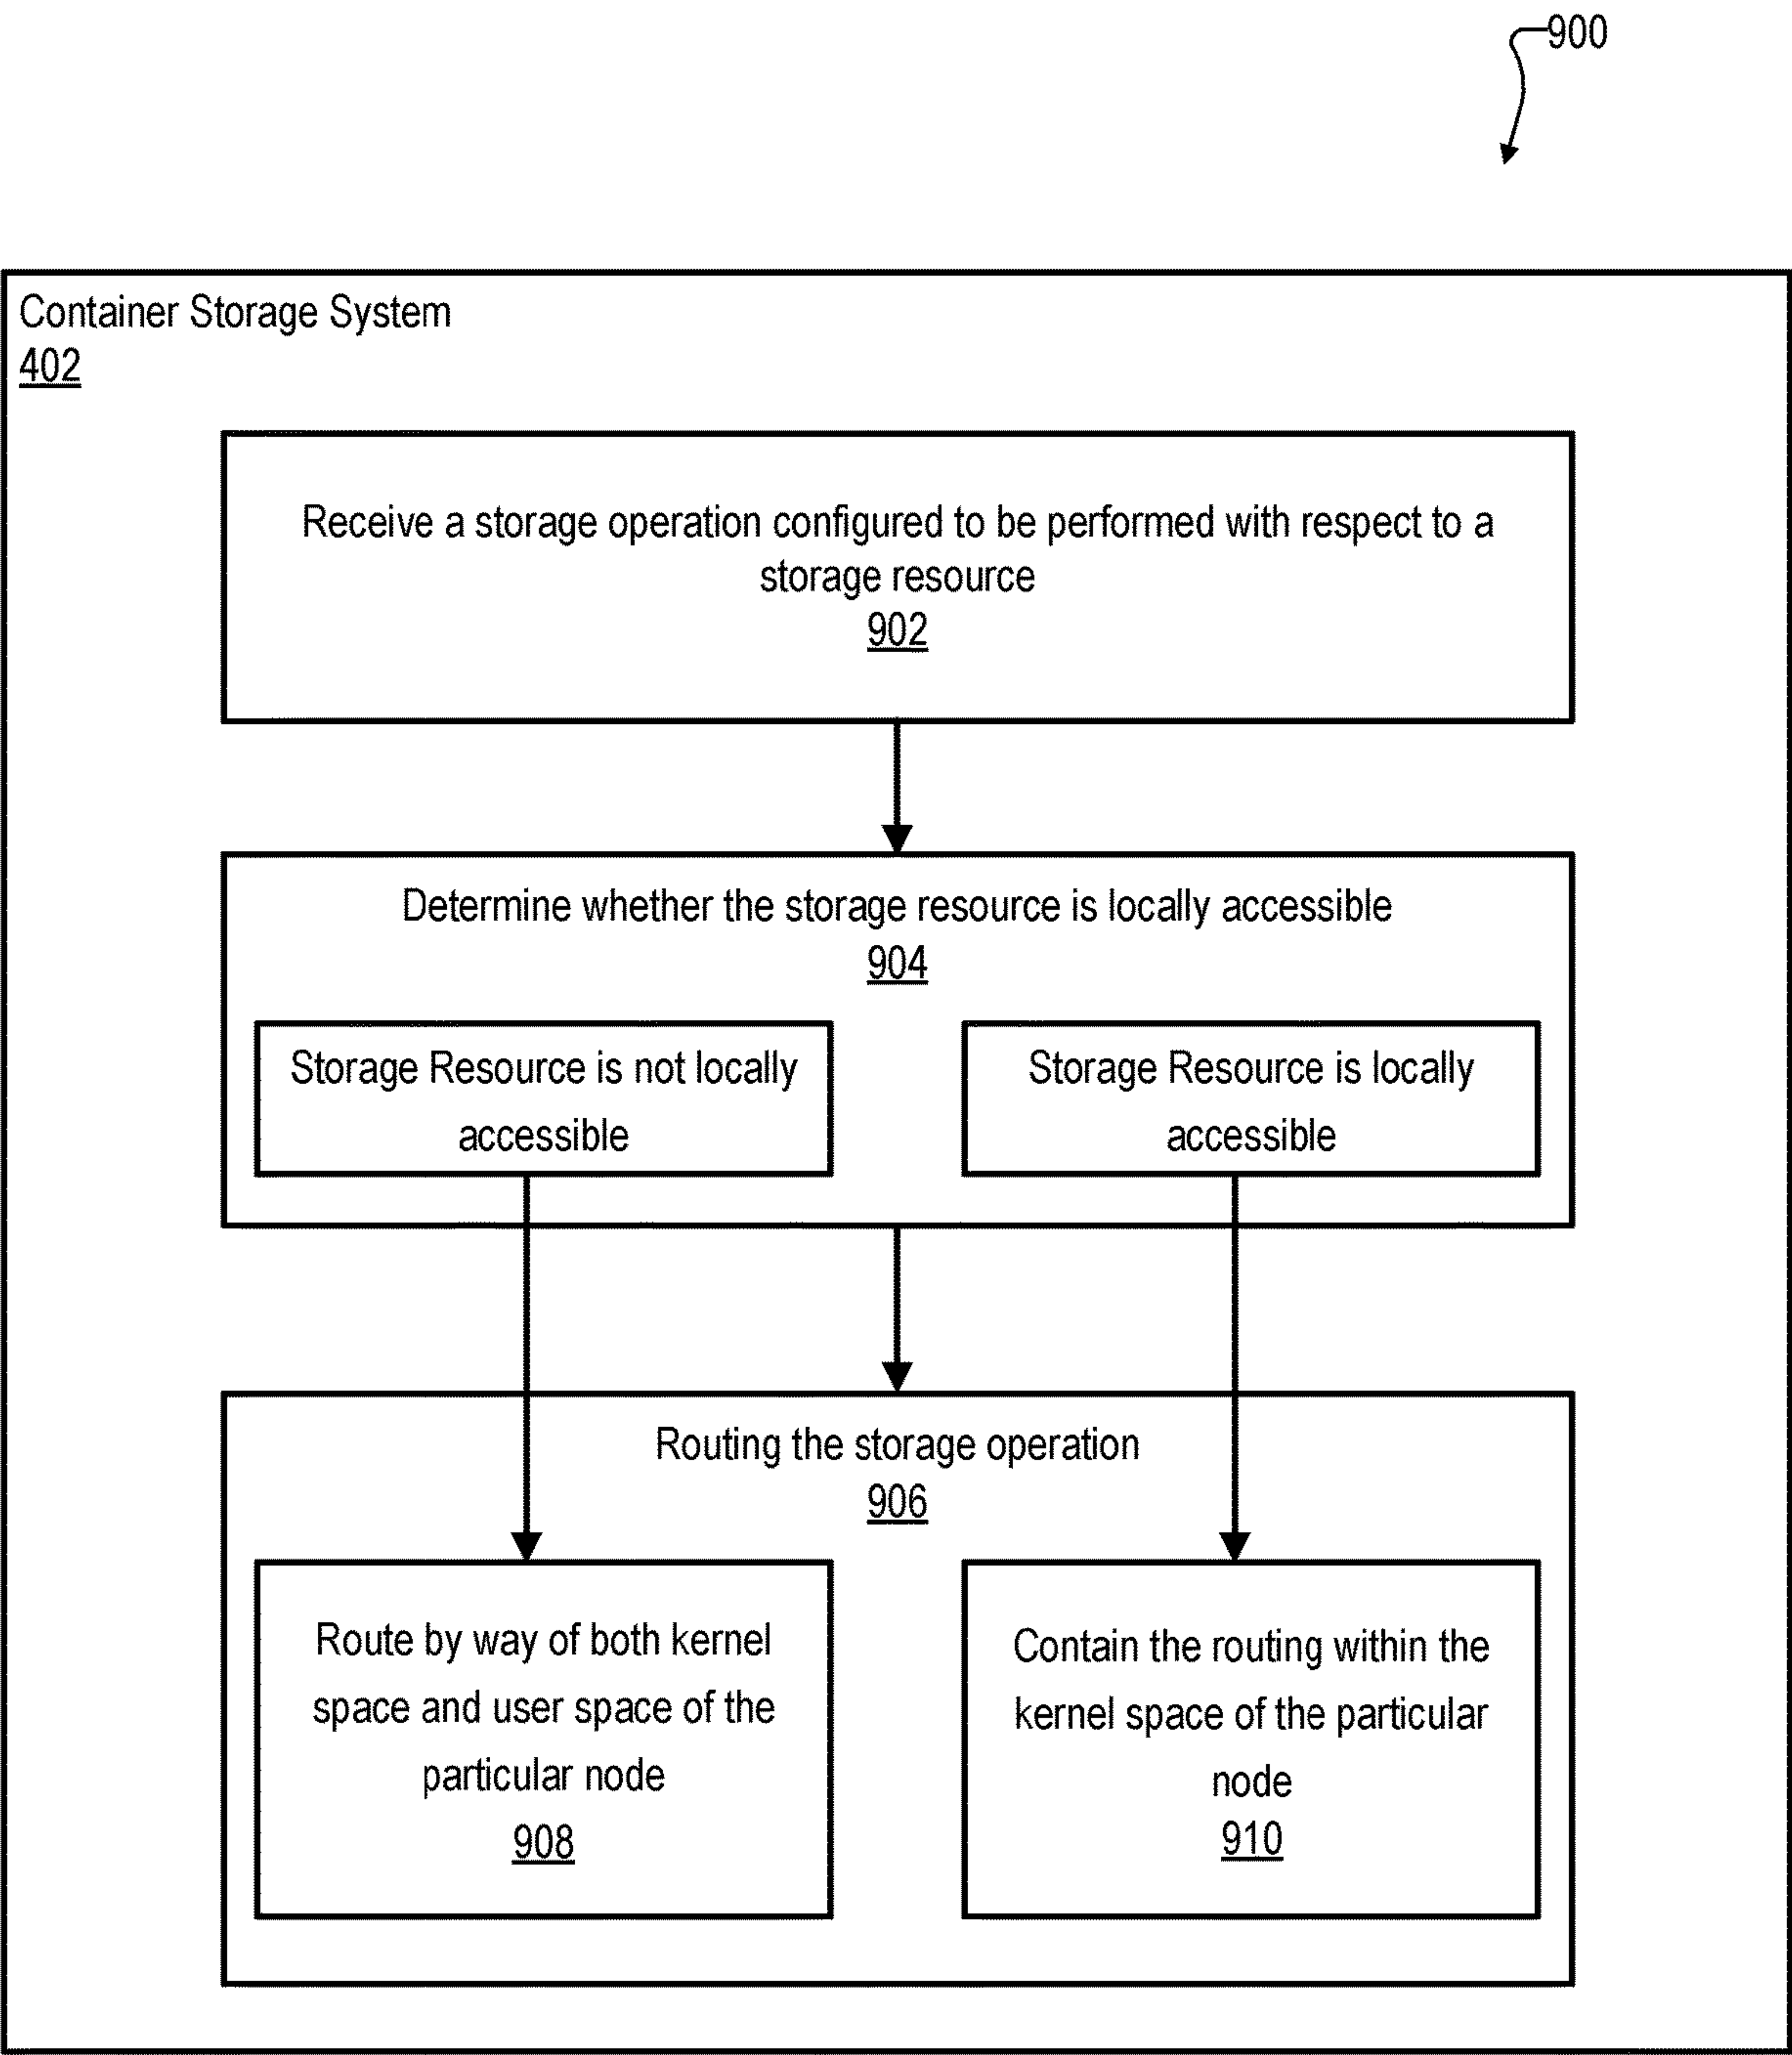


FIG. 9




FIG. 10




FIG. 11




FIG. 12




FIG. 13

STORAGE OPERATION ROUTING IN A CONTAINER SYSTEM

RELATED APPLICATIONS

[0001] The present application is a continuation-in-part application of U.S. patent application Ser. No. 17/514,872, filed Oct. 29, 2021 and incorporated by reference herein in its entirety.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] The accompanying drawings illustrate various embodiments and are a part of the specification. The illustrated embodiments are merely examples and do not limit the scope of the disclosure. Throughout the drawings, identical or similar reference numbers designate identical or similar elements.

[0003] FIG. 1A illustrates a first example system for data storage in accordance with some implementations.

[0004] FIG. 1B illustrates a second example system for data storage in accordance with some implementations.

[0005] FIG. 1C illustrates a third example system for data storage in accordance with some implementations.

[0006] FIG. 1D illustrates a fourth example system for data storage in accordance with some implementations.

[0007] FIG. 2A is a perspective view of a storage cluster with multiple storage nodes and internal storage coupled to each storage node to provide network attached storage, in accordance with some embodiments.

[0008] FIG. 2B is a block diagram showing an interconnect switch coupling multiple storage nodes in accordance with some embodiments.

[0009] FIG. 2C is a multiple level block diagram, showing contents of a storage node and contents of one of the non-volatile solid state storage units in accordance with some embodiments.

[0010] FIG. 2D shows a storage server environment, which uses embodiments of the storage nodes and storage units of some previous figures in accordance with some embodiments.

[0011] FIG. 2E is a blade hardware block diagram, showing a control plane, compute and storage planes, and authorities interacting with underlying physical resources, in accordance with some embodiments.

[0012] FIG. 2F depicts elasticity software layers in blades of a storage cluster, in accordance with some embodiments.

[0013] FIG. 2G depicts authorities and storage resources in blades of a storage cluster, in accordance with some embodiments.

[0014] FIG. 3A sets forth a diagram of a storage system that is coupled for data communications with a cloud services provider in accordance with some embodiments of the present disclosure.

[0015] FIG. 3B sets forth a diagram of a storage system in accordance with some embodiments of the present disclosure.

[0016] FIG. 3C sets forth an example of a cloud-based storage system in accordance with some embodiments of the present disclosure.

[0017] FIG. 3D illustrates an example computing device that may be specifically configured to perform one or more of the processes described herein.

[0018] FIG. 3E illustrates an example of a fleet of storage systems for providing storage services in accordance with some embodiments of the present disclosure.

[0019] FIGS. 4A and 4B illustrate example container storage systems in accordance with some embodiments of the present disclosure.

[0020] FIG. 5 illustrates storage path routing within a container system, in accordance with some embodiments of the present disclosure.

[0021] FIG. 6 depicts an example method of a storage path routing in a container system, in accordance with some embodiments of the present disclosure.

[0022] FIG. 7 depicts an example method of a storage path routing in a container system, in accordance with embodiments of the present disclosure.

[0023] FIG. 8 depicts an example method of a storage path routing in a container system, in accordance with some embodiments of the present disclosure.

[0024] FIG. 9 depicts an example method of routing a storage operation in a container system in accordance with some embodiments of the present disclosure.

[0025] FIG. 10 illustrates a configuration in which various routes through a kernel space and/or a user space of a node may be implemented for routing a storage operation to a storage resource in accordance with some embodiments of the present disclosure.

[0026] FIG. 11 show illustrative implementations for routing storage operations that are configured to be performed with respect to multiple storage resources in accordance with some embodiments of the present disclosure.

[0027] FIG. 12 illustrates a configuration in which various routes through a kernel space and/or a user space of a node may be implemented for routing a storage operation to a storage resource in accordance with some embodiments of the present disclosure.

[0028] FIG. 13 illustrates a configuration in which various routes through a kernel space and/or a user space of a node may be implemented for routing a storage operation to a storage resource in accordance with some embodiments of the present disclosure.

DESCRIPTION OF EMBODIMENTS

[0029] Example methods, apparatuses, and products for storage operation routing in a container system in accordance with embodiments of the present disclosure are described with reference to the accompanying drawings, beginning with FIG. 1A. FIG. 1A illustrates an example system for data storage, in accordance with some implementations. System 100 (also referred to as “storage system” herein) includes numerous elements for purposes of illustration rather than limitation. It may be noted that system 100 may include the same, more, or fewer elements configured in the same or different manner in other implementations.

[0030] System 100 includes a number of computing devices 164A-B. Computing devices (also referred to as “client devices” herein) may be embodied, for example, a server in a data center, a workstation, a personal computer, a notebook, or the like. Computing devices 164A-B may be coupled for data communications to one or more storage arrays 102A-B through a storage area network (‘SAN’) 158 or a local area network (‘LAN’) 160.

[0031] The SAN 158 may be implemented with a variety of data communications fabrics, devices, and protocols. For

example, the fabrics for SAN 158 may include Fibre Channel, Ethernet, Infiniband, Serial Attached Small Computer System Interface ('SAS'), or the like. Data communications protocols for use with SAN 158 may include Advanced Technology Attachment ('ATA'), Fibre Channel Protocol, Small Computer System Interface ('SCSI'), Internet Small Computer System Interface ('iSCSI'), HyperSCSI, Non-Volatile Memory Express ('NVMe') over Fabrics, or the like. It may be noted that SAN 158 is provided for illustration, rather than limitation. Other data communication couplings may be implemented between computing devices 164A-B and storage arrays 102A-B.

[0032] The LAN 160 may also be implemented with a variety of fabrics, devices, and protocols. For example, the fabrics for LAN 160 may include Ethernet (802.3), wireless (802.11), or the like. Data communication protocols for use in LAN 160 may include Transmission Control Protocol ('TCP'), User Datagram Protocol ('UDP'), Internet Protocol ('IP'), HyperText Transfer Protocol ('HTTP'), Wireless Access Protocol ('WAP'), Handheld Device Transport Protocol ('HDTP'), Session Initiation Protocol ('SIP'), Real Time Protocol ('RTP'), or the like.

[0033] Storage arrays 102A-B may provide persistent data storage for the computing devices 164A-B. Storage array 102A may be contained in a chassis (not shown), and storage array 102B may be contained in another chassis (not shown), in implementations. Storage array 102A and 102B may include one or more storage array controllers 110A-D (also referred to as "controller" herein). A storage array controller 110A-D may be embodied as a module of automated computing machinery comprising computer hardware, computer software, or a combination of computer hardware and software. In some implementations, the storage array controllers 110A-D may be configured to carry out various storage tasks. Storage tasks may include writing data received from the computing devices 164A-B to storage array 102A-B, erasing data from storage array 102A-B, retrieving data from storage array 102A-B and providing data to computing devices 164A-B, monitoring and reporting of disk utilization and performance, performing redundancy operations, such as Redundant Array of Independent Drives ('RAID') or RAID-like data redundancy operations, compressing data, encrypting data, and so forth.

[0034] Storage array controller 110A-D may be implemented in a variety of ways, including as a Field Programmable Gate Array ('FPGA'), a Programmable Logic Chip ('PLC'), an Application Specific Integrated Circuit ('ASIC'), System-on-Chip ('SOC'), or any computing device that includes discrete components such as a processing device, central processing unit, computer memory, or various adapters. Storage array controller 110A-D may include, for example, a data communications adapter configured to support communications via the SAN 158 or LAN 160. In some implementations, storage array controller 110A-D may be independently coupled to the LAN 160. In implementations, storage array controller 110A-D may include an I/O controller or the like that couples the storage array controller 110A-D for data communications, through a midplane (not shown), to a persistent storage resource 170A-B (also referred to as a "storage resource" herein). The persistent storage resource 170A-B may include any number of storage drives 171A-F (also referred to as "storage devices" herein) and any number of non-volatile Random Access Memory ('NVRAM') devices (not shown).

[0035] In some implementations, the NVRAM devices of a persistent storage resource 170A-B may be configured to receive, from the storage array controller 110A-D, data to be stored in the storage drives 171A-F. In some examples, the data may originate from computing devices 164A-B. In some examples, writing data to the NVRAM device may be carried out more quickly than directly writing data to the storage drive 171A-F. In implementations, the storage array controller 110A-D may be configured to utilize the NVRAM devices as a quickly accessible buffer for data destined to be written to the storage drives 171A-F. Latency for write requests using NVRAM devices as a buffer may be improved relative to a system in which a storage array controller 110A-D writes data directly to the storage drives 171A-F. In some implementations, the NVRAM devices may be implemented with computer memory in the form of high bandwidth, low latency RAM. The NVRAM device is referred to as "non-volatile" because the NVRAM device may receive or include a unique power source that maintains the state of the RAM after main power loss to the NVRAM device. Such a power source may be a battery, one or more capacitors, or the like. In response to a power loss, the NVRAM device may be configured to write the contents of the RAM to a persistent storage, such as the storage drives 171A-F.

[0036] In implementations, storage drive 171A-F may refer to any device configured to record data persistently, where "persistently" or "persistent" refers as to a device's ability to maintain recorded data after loss of power. In some implementations, storage drive 171A-F may correspond to non-disk storage media. For example, the storage drive 171A-F may be one or more solid-state drives ('SSDs'), flash memory based storage, any type of solid-state non-volatile memory, or any other type of non-mechanical storage device. In other implementations, storage drive 171A-F may include mechanical or spinning hard disk, such as hard-disk drives ('HDD').

[0037] In some implementations, the storage array controllers 110A-D may be configured for offloading device management responsibilities from storage drive 171A-F in storage array 102A-B. For example, storage array controllers 110A-D may manage control information that may describe the state of one or more memory blocks in the storage drives 171A-F. The control information may indicate, for example, that a particular memory block has failed and should no longer be written to, that a particular memory block contains boot code for a storage array controller 110A-D, the number of program-erase ('P/E') cycles that have been performed on a particular memory block, the age of data stored in a particular memory block, the type of data that is stored in a particular memory block, and so forth. In some implementations, the control information may be stored with an associated memory block as metadata. In other implementations, the control information for the storage drives 171A-F may be stored in one or more particular memory blocks of the storage drives 171A-F that are selected by the storage array controller 110A-D. The selected memory blocks may be tagged with an identifier indicating that the selected memory block contains control information. The identifier may be utilized by the storage array controllers 110A-D in conjunction with storage drives 171A-F to quickly identify the memory blocks that contain control information. For example, the storage controllers 110A-D may issue a command to locate memory blocks that

contain control information. It may be noted that control information may be so large that parts of the control information may be stored in multiple locations, that the control information may be stored in multiple locations for purposes of redundancy, for example, or that the control information may otherwise be distributed across multiple memory blocks in the storage drive 171A-F.

[0038] In implementations, storage array controllers 110A-D may offload device management responsibilities from storage drives 171A-F of storage array 102A-B by retrieving, from the storage drives 171A-F, control information describing the state of one or more memory blocks in the storage drives 171A-F. Retrieving the control information from the storage drives 171A-F may be carried out, for example, by the storage array controller 110A-D querying the storage drives 171A-F for the location of control information for a particular storage drive 171A-F. The storage drives 171A-F may be configured to execute instructions that enable the storage drive 171A-F to identify the location of the control information. The instructions may be executed by a controller (not shown) associated with or otherwise located on the storage drive 171A-F and may cause the storage drive 171A-F to scan a portion of each memory block to identify the memory blocks that store control information for the storage drives 171A-F. The storage drives 171A-F may respond by sending a response message to the storage array controller 110A-D that includes the location of control information for the storage drive 171A-F. Responsive to receiving the response message, storage array controllers 110A-D may issue a request to read data stored at the address associated with the location of control information for the storage drives 171A-F.

[0039] In other implementations, the storage array controllers 110A-D may further offload device management responsibilities from storage drives 171A-F by performing, in response to receiving the control information, a storage drive management operation. A storage drive management operation may include, for example, an operation that is typically performed by the storage drive 171A-F (e.g., the controller (not shown) associated with a particular storage drive 171A-F). A storage drive management operation may include, for example, ensuring that data is not written to failed memory blocks within the storage drive 171A-F, ensuring that data is written to memory blocks within the storage drive 171A-F in such a way that adequate wear leveling is achieved, and so forth.

[0040] In implementations, storage array 102A-B may implement two or more storage array controllers 110A-D. For example, storage array 102A may include storage array controllers 110A and storage array controllers 110B. At a given instance, a single storage array controller 110A-D (e.g., storage array controller 110A) of a storage system 100 may be designated with primary status (also referred to as “primary controller” herein), and other storage array controllers 110A-D (e.g., storage array controller 110B) may be designated with secondary status (also referred to as “secondary controller” herein). The primary controller may have particular rights, such as permission to alter data in persistent storage resource 170A-B (e.g., writing data to persistent storage resource 170A-B). At least some of the rights of the primary controller may supersede the rights of the secondary controller. For instance, the secondary controller may not have permission to alter data in persistent storage resource 170A-B when the primary controller has the right. The status

of storage array controllers 110A-D may change. For example, storage array controller 110A may be designated with secondary status, and storage array controller 110B may be designated with primary status.

[0041] In some implementations, a primary controller, such as storage array controller 110A, may serve as the primary controller for one or more storage arrays 102A-B, and a second controller, such as storage array controller 110B, may serve as the secondary controller for the one or more storage arrays 102A-B. For example, storage array controller 110A may be the primary controller for storage array 102A and storage array 102B, and storage array controller 110B may be the secondary controller for storage array 102A and 102B. In some implementations, storage array controllers 110C and 110D (also referred to as “storage processing modules”) may neither have primary or secondary status. Storage array controllers 110C and 110D, implemented as storage processing modules, may act as a communication interface between the primary and secondary controllers (e.g., storage array controllers 110A and 110B, respectively) and storage array 102B. For example, storage array controller 110A of storage array 102A may send a write request, via SAN 158, to storage array 102B. The write request may be received by both storage array controllers 110C and 110D of storage array 102B. Storage array controllers 110C and 110D facilitate the communication, e.g., send the write request to the appropriate storage drive 171A-F. It may be noted that in some implementations storage processing modules may be used to increase the number of storage drives controlled by the primary and secondary controllers.

[0042] In implementations, storage array controllers 110A-D are communicatively coupled, via a midplane (not shown), to one or more storage drives 171A-F and to one or more NVRAM devices (not shown) that are included as part of a storage array 102A-B. The storage array controllers 110A-D may be coupled to the midplane via one or more data communication links and the midplane may be coupled to the storage drives 171A-F and the NVRAM devices via one or more data communications links. The data communications links described herein are collectively illustrated by data communications links 108A-D and may include a Peripheral Component Interconnect Express (‘PCIe’) bus, for example.

[0043] FIG. 1B illustrates an example system for data storage, in accordance with some implementations. Storage array controller 101 illustrated in FIG. 1B may be similar to the storage array controllers 110A-D described with respect to FIG. 1A. In one example, storage array controller 101 may be similar to storage array controller 110A or storage array controller 110B. Storage array controller 101 includes numerous elements for purposes of illustration rather than limitation. It may be noted that storage array controller 101 may include the same, more, or fewer elements configured in the same or different manner in other implementations. It may be noted that elements of FIG. 1A may be included below to help illustrate features of storage array controller 101.

[0044] Storage array controller 101 may include one or more processing devices 104 and random access memory (‘RAM’) 111. Processing device 104 (or controller 101) represents one or more general-purpose processing devices such as a microprocessor, central processing unit, or the like. More particularly, the processing device 104 (or controller

101) may be a complex instruction set computing ('CISC') microprocessor, reduced instruction set computing ('RISC') microprocessor, very long instruction word ('VLIW') microprocessor, or a processor implementing other instruction sets or processors implementing a combination of instruction sets. The processing device **104** (or controller **101**) may also be one or more special-purpose processing devices such as an ASIC, an FPGA, a digital signal processor ('DSP'), network processor, or the like.

[0045] The processing device **104** may be connected to the RAM **111** via a data communications link **106**, which may be embodied as a high speed memory bus such as a Double-Data Rate 4 ('DDR4') bus. Stored in RAM **111** is an operating system **112**. In some implementations, instructions **113** are stored in RAM **111**. Instructions **113** may include computer program instructions for performing operations in a direct-mapped flash storage system. In one embodiment, a direct-mapped flash storage system is one that addresses data blocks within flash drives directly and without an address translation performed by the storage controllers of the flash drives.

[0046] In implementations, storage array controller **101** includes one or more host bus adapters **103A-C** that are coupled to the processing device **104** via a data communications link **105A-C**. In implementations, host bus adapters **103A-C** may be computer hardware that connects a host system (e.g., the storage array controller) to other network and storage arrays. In some examples, host bus adapters **103A-C** may be a Fibre Channel adapter that enables the storage array controller **101** to connect to a SAN, an Ethernet adapter that enables the storage array controller **101** to connect to a LAN, or the like. Host bus adapters **103A-C** may be coupled to the processing device **104** via a data communications link **105A-C** such as, for example, a PCIe bus.

[0047] In implementations, storage array controller **101** may include a host bus adapter **114** that is coupled to an expander **115**. The expander **115** may be used to attach a host system to a larger number of storage drives. The expander **115** may, for example, be a SAS expander utilized to enable the host bus adapter **114** to attach to storage drives in an implementation where the host bus adapter **114** is embodied as a SAS controller.

[0048] In implementations, storage array controller **101** may include a switch **116** coupled to the processing device **104** via a data communications link **109**. The switch **116** may be a computer hardware device that can create multiple endpoints out of a single endpoint, thereby enabling multiple devices to share a single endpoint. The switch **116** may, for example, be a PCIe switch that is coupled to a PCIe bus (e.g., data communications link **109**) and presents multiple PCIe connection points to the midplane.

[0049] In implementations, storage array controller **101** includes a data communications link **107** for coupling the storage array controller **101** to other storage array controllers. In some examples, data communications link **107** may be a QuickPath Interconnect (QPI) interconnect.

[0050] A traditional storage system that uses traditional flash drives may implement a process across the flash drives that are part of the traditional storage system. For example, a higher level process of the storage system may initiate and control a process across the flash drives. However, a flash drive of the traditional storage system may include its own storage controller that also performs the process. Thus, for

the traditional storage system, a higher level process (e.g., initiated by the storage system) and a lower level process (e.g., initiated by a storage controller of the storage system) may both be performed.

[0051] To resolve various deficiencies of a traditional storage system, operations may be performed by higher level processes and not by the lower level processes. For example, the flash storage system may include flash drives that do not include storage controllers that provide the process. Thus, the operating system of the flash storage system itself may initiate and control the process. This may be accomplished by a direct-mapped flash storage system that addresses data blocks within the flash drives directly and without an address translation performed by the storage controllers of the flash drives.

[0052] In implementations, storage drive **171A-F** may be one or more zoned storage devices. In some implementations, the one or more zoned storage devices may be a shingled HDD. In implementations, the one or more storage devices may be a flash-based SSD. In a zoned storage device, a zoned namespace on the zoned storage device can be addressed by groups of blocks that are grouped and aligned by a natural size, forming a number of addressable zones. In implementations utilizing an SSD, the natural size may be based on the erase block size of the SSD. In some implementations, the zones of the zoned storage device may be defined during initialization of the zoned storage device. In implementations, the zones may be defined dynamically as data is written to the zoned storage device.

[0053] In some implementations, zones may be heterogeneous, with some zones each being a page group and other zones being multiple page groups. In implementations, some zones may correspond to an erase block and other zones may correspond to multiple erase blocks. In an implementation, zones may be any combination of differing numbers of pages in page groups and/or erase blocks, for heterogeneous mixes of programming modes, manufacturers, product types and/or product generations of storage devices, as applied to heterogeneous assemblies, upgrades, distributed storages, etc. In some implementations, zones may be defined as having usage characteristics, such as a property of supporting data with particular kinds of longevity (very short lived or very long lived, for example). These properties could be used by a zoned storage device to determine how the zone will be managed over the zone's expected lifetime.

[0054] It should be appreciated that a zone is a virtual construct. Any particular zone may not have a fixed location at a storage device. Until allocated, a zone may not have any location at a storage device. A zone may correspond to a number representing a chunk of virtually allocatable space that is the size of an erase block or other block size in various implementations. When the system allocates or opens a zone, zones get allocated to flash or other solid-state storage memory and, as the system writes to the zone, pages are written to that mapped flash or other solid-state storage memory of the zoned storage device. When the system closes the zone, the associated erase block(s) or other sized block(s) are completed. At some point in the future, the system may delete a zone which will free up the zone's allocated space. During its lifetime, a zone may be moved around to different locations of the zoned storage device, e.g., as the zoned storage device does internal maintenance.

[0055] In implementations, the zones of the zoned storage device may be in different states. A zone may be in an empty

state in which data has not been stored at the zone. An empty zone may be opened explicitly, or implicitly by writing data to the zone. This is the initial state for zones on a fresh zoned storage device, but may also be the result of a zone reset. In some implementations, an empty zone may have a designated location within the flash memory of the zoned storage device. In an implementation, the location of the empty zone may be chosen when the zone is first opened or first written to (or later if writes are buffered into memory). A zone may be in an open state either implicitly or explicitly, where a zone that is in an open state may be written to store data with write or append commands. In an implementation, a zone that is in an open state may also be written to using a copy command that copies data from a different zone. In some implementations, a zoned storage device may have a limit on the number of open zones at a particular time.

[0056] A zone in a closed state is a zone that has been partially written to, but has entered a closed state after issuing an explicit close operation. A zone in a closed state may be left available for future writes, but may reduce some of the run-time overhead consumed by keeping the zone in an open state. In implementations, a zoned storage device may have a limit on the number of closed zones at a particular time. A zone in a full state is a zone that is storing data and can no longer be written to. A zone may be in a full state either after writes have written data to the entirety of the zone or as a result of a zone finish operation. Prior to a finish operation, a zone may or may not have been completely written. After a finish operation, however, the zone may not be opened or written to further without first performing a zone reset operation.

[0057] The mapping from a zone to an erase block (or to a shingled track in an HDD) may be arbitrary, dynamic, and hidden from view. The process of opening a zone may be an operation that allows a new zone to be dynamically mapped to underlying storage of the zoned storage device, and then allows data to be written through appending writes into the zone until the zone reaches capacity. The zone can be finished at any point, after which further data may not be written into the zone. When the data stored at the zone is no longer needed, the zone can be reset which effectively deletes the zone's content from the zoned storage device, making the physical storage held by that zone available for the subsequent storage of data. Once a zone has been written and finished, the zoned storage device ensures that the data stored at the zone is not lost until the zone is reset. In the time between writing the data to the zone and the resetting of the zone, the zone may be moved around between shingle tracks or erase blocks as part of maintenance operations within the zoned storage device, such as by copying data to keep the data refreshed or to handle memory cell aging in an SSD.

[0058] In implementations utilizing an HDD, the resetting of the zone may allow the shingle tracks to be allocated to a new, opened zone that may be opened at some point in the future. In implementations utilizing an SSD, the resetting of the zone may cause the associated physical erase block(s) of the zone to be erased and subsequently reused for the storage of data. In some implementations, the zoned storage device may have a limit on the number of open zones at a point in time to reduce the amount of overhead dedicated to keeping zones open.

[0059] The operating system of the flash storage system may identify and maintain a list of allocation units across

multiple flash drives of the flash storage system. The allocation units may be entire erase blocks or multiple erase blocks. The operating system may maintain a map or address range that directly maps addresses to erase blocks of the flash drives of the flash storage system.

[0060] Direct mapping to the erase blocks of the flash drives may be used to rewrite data and erase data. For example, the operations may be performed on one or more allocation units that include a first data and a second data where the first data is to be retained and the second data is no longer being used by the flash storage system. The operating system may initiate the process to write the first data to new locations within other allocation units and erasing the second data and marking the allocation units as being available for use for subsequent data. Thus, the process may only be performed by the higher level operating system of the flash storage system without an additional lower level process being performed by controllers of the flash drives.

[0061] Advantages of the process being performed only by the operating system of the flash storage system include increased reliability of the flash drives of the flash storage system as unnecessary or redundant write operations are not being performed during the process. One possible point of novelty here is the concept of initiating and controlling the process at the operating system of the flash storage system. In addition, the process can be controlled by the operating system across multiple flash drives. This is contrast to the process being performed by a storage controller of a flash drive.

[0062] A storage system can consist of two storage array controllers that share a set of drives for failover purposes, or it could consist of a single storage array controller that provides a storage service that utilizes multiple drives, or it could consist of a distributed network of storage array controllers each with some number of drives or some amount of Flash storage where the storage array controllers in the network collaborate to provide a complete storage service and collaborate on various aspects of a storage service including storage allocation and garbage collection.

[0063] FIG. 1C illustrates a third example system 117 for data storage in accordance with some implementations. System 117 (also referred to as "storage system" herein) includes numerous elements for purposes of illustration rather than limitation. It may be noted that system 117 may include the same, more, or fewer elements configured in the same or different manner in other implementations.

[0064] In one embodiment, system 117 includes a dual Peripheral Component Interconnect (PCI) flash storage device 118 with separately addressable fast write storage. System 117 may include a storage device controller 119. In one embodiment, storage device controller 119A-D may be a CPU, ASIC, FPGA, or any other circuitry that may implement control structures necessary according to the present disclosure. In one embodiment, system 117 includes flash memory devices (e.g., including flash memory devices 120a-n), operatively coupled to various channels of the storage device controller 119. Flash memory devices 120a-n, may be presented to the controller 119A-D as an addressable collection of Flash pages, erase blocks, and/or control elements sufficient to allow the storage device controller 119A-D to program and retrieve various aspects of the Flash. In one embodiment, storage device controller 119A-D may perform operations on flash memory devices 120a-n includ-

ing storing and retrieving data content of pages, arranging and erasing any blocks, tracking statistics related to the use and reuse of Flash memory pages, erase blocks, and cells, tracking and predicting error codes and faults within the Flash memory, controlling voltage levels associated with programming and retrieving contents of Flash cells, etc.

[0065] In one embodiment, system **117** may include RAM **121** to store separately addressable fast-write data. In one embodiment, RAM **121** may be one or more separate discrete devices. In another embodiment, RAM **121** may be integrated into storage device controller **119A-D** or multiple storage device controllers. The RAM **121** may be utilized for other purposes as well, such as temporary program memory for a processing device (e.g., a CPU) in the storage device controller **119**.

[0066] In one embodiment, system **117** may include a stored energy device **122**, such as a rechargeable battery or a capacitor. Stored energy device **122** may store energy sufficient to power the storage device controller **119**, some amount of the RAM (e.g., RAM **121**), and some amount of Flash memory (e.g., Flash memory **120a-120n**) for sufficient time to write the contents of RAM to Flash memory. In one embodiment, storage device controller **119A-D** may write the contents of RAM to Flash Memory if the storage device controller detects loss of external power.

[0067] In one embodiment, system **117** includes two data communications links **123a**, **123b**. In one embodiment, data communications links **123a**, **123b** may be PCI interfaces. In another embodiment, data communications links **123a**, **123b** may be based on other communications standards (e.g., HyperTransport, InfiniBand, etc.). Data communications links **123a**, **123b** may be based on non-volatile memory express ('NVMe') or NVMe over fabrics ('NVMeF') specifications that allow external connection to the storage device controller **119A-D** from other components in the storage system **117**. It should be noted that data communications links may be interchangeably referred to herein as PCI buses for convenience.

[0068] System **117** may also include an external power source (not shown), which may be provided over one or both data communications links **123a**, **123b**, or which may be provided separately. An alternative embodiment includes a separate Flash memory (not shown) dedicated for use in storing the content of RAM **121**. The storage device controller **119A-D** may present a logical device over a PCI bus which may include an addressable fast-write logical device, or a distinct part of the logical address space of the storage device **118**, which may be presented as PCI memory or as persistent storage. In one embodiment, operations to store into the device are directed into the RAM **121**. On power failure, the storage device controller **119A-D** may write stored content associated with the addressable fast-write logical storage to Flash memory (e.g., Flash memory **120a-n**) for long-term persistent storage.

[0069] In one embodiment, the logical device may include some presentation of some or all of the content of the Flash memory devices **120a-n**, where that presentation allows a storage system including a storage device **118** (e.g., storage system **117**) to directly address Flash memory pages and directly reprogram erase blocks from storage system components that are external to the storage device through the PCI bus. The presentation may also allow one or more of the external components to control and retrieve other aspects of the Flash memory including some or all of: tracking statis-

tics related to use and reuse of Flash memory pages, erase blocks, and cells across all the Flash memory devices; tracking and predicting error codes and faults within and across the Flash memory devices; controlling voltage levels associated with programming and retrieving contents of Flash cells; etc.

[0070] In one embodiment, the stored energy device **122** may be sufficient to ensure completion of in-progress operations to the Flash memory devices **120a-120n** stored energy device **122** may power storage device controller **119A-D** and associated Flash memory devices (e.g., **120a-n**) for those operations, as well as for the storing of fast-write RAM to Flash memory. Stored energy device **122** may be used to store accumulated statistics and other parameters kept and tracked by the Flash memory devices **120a-n** and/or the storage device controller **119**. Separate capacitors or stored energy devices (such as smaller capacitors near or embedded within the Flash memory devices themselves) may be used for some or all of the operations described herein.

[0071] Various schemes may be used to track and optimize the life span of the stored energy component, such as adjusting voltage levels over time, partially discharging the stored energy device **122** to measure corresponding discharge characteristics, etc. If the available energy decreases over time, the effective available capacity of the addressable fast-write storage may be decreased to ensure that it can be written safely based on the currently available stored energy.

[0072] FIG. 1D illustrates a third example storage system **124** for data storage in accordance with some implementations. In one embodiment, storage system **124** includes storage controllers **125a**, **125b**. In one embodiment, storage controllers **125a**, **125b** are operatively coupled to Dual PCI storage devices. Storage controllers **125a**, **125b** may be operatively coupled (e.g., via a storage network **130**) to some number of host computers **127a-n**.

[0073] In one embodiment, two storage controllers (e.g., **125a** and **125b**) provide storage services, such as a SCS) block storage array, a file server, an object server, a database or data analytics service, etc. The storage controllers **125a**, **125b** may provide services through some number of network interfaces (e.g., **126a-d**) to host computers **127a-n** outside of the storage system **124**. Storage controllers **125a**, **125b** may provide integrated services or an application entirely within the storage system **124**, forming a converged storage and compute system. The storage controllers **125a**, **125b** may utilize the fast write memory within or across storage devices **119a-d** to journal in progress operations to ensure the operations are not lost on a power failure, storage controller removal, storage controller or storage system shutdown, or some fault of one or more software or hardware components within the storage system **124**.

[0074] In one embodiment, storage controllers **125a**, **125b** operate as PCI masters to one or the other PCI buses **128a**, **128b**. In another embodiment, **128a** and **128b** may be based on other communications standards (e.g., HyperTransport, InfiniBand, etc.). Other storage system embodiments may operate storage controllers **125a**, **125b** as multi-masters for both PCI buses **128a**, **128b**. Alternately, a PCI/NVMe/NVMeF switching infrastructure or fabric may connect multiple storage controllers. Some storage system embodiments may allow storage devices to communicate with each other directly rather than communicating only with storage controllers. In one embodiment, a storage device controller **119a** may be operable under direction from a storage controller

125a to synthesize and transfer data to be stored into Flash memory devices from data that has been stored in RAM (e.g., RAM **121** of FIG. 1C). For example, a recalculated version of RAM content may be transferred after a storage controller has determined that an operation has fully committed across the storage system, or when fast-write memory on the device has reached a certain used capacity, or after a certain amount of time, to ensure improve safety of the data or to release addressable fast-write capacity for reuse. This mechanism may be used, for example, to avoid a second transfer over a bus (e.g., **128a**, **128b**) from the storage controllers **125a**, **125b**. In one embodiment, a recalculation may include compressing data, attaching indexing or other metadata, combining multiple data segments together, performing erasure code calculations, etc.

[0075] In one embodiment, under direction from a storage controller **125a**, **125b**, a storage device controller **119a**, **119b** may be operable to calculate and transfer data to other storage devices from data stored in RAM (e.g., RAM **121** of FIG. 1C) without involvement of the storage controllers **125a**, **125b**. This operation may be used to mirror data stored in one storage controller **125a** to another storage controller **125b**, or it could be used to offload compression, data aggregation, and/or erasure coding calculations and transfers to storage devices to reduce load on storage controllers or the storage controller interface **129a**, **129b** to the PCI bus **128a**, **128b**.

[0076] A storage device controller **119A-D** may include mechanisms for implementing high availability primitives for use by other parts of a storage system external to the Dual PCI storage device **118**. For example, reservation or exclusion primitives may be provided so that, in a storage system with two storage controllers providing a highly available storage service, one storage controller may prevent the other storage controller from accessing or continuing to access the storage device. This could be used, for example, in cases where one controller detects that the other controller is not functioning properly or where the interconnect between the two storage controllers may itself not be functioning properly.

[0077] In one embodiment, a storage system for use with Dual PCI direct mapped storage devices with separately addressable fast write storage includes systems that manage erase blocks or groups of erase blocks as allocation units for storing data on behalf of the storage service, or for storing metadata (e.g., indexes, logs, etc.) associated with the storage service, or for proper management of the storage system itself. Flash pages, which may be a few kilobytes in size, may be written as data arrives or as the storage system is to persist data for long intervals of time (e.g., above a defined threshold of time). To commit data more quickly, or to reduce the number of writes to the Flash memory devices, the storage controllers may first write data into the separately addressable fast write storage on one more storage devices.

[0078] In one embodiment, the storage controllers **125a**, **125b** may initiate the use of erase blocks within and across storage devices (e.g., **118**) in accordance with an age and expected remaining lifespan of the storage devices, or based on other statistics. The storage controllers **125a**, **125b** may initiate garbage collection and data migration data between storage devices in accordance with pages that are no longer needed as well as to manage Flash page and erase block lifespans and to manage overall system performance.

[0079] In one embodiment, the storage system **124** may utilize mirroring and/or erasure coding schemes as part of storing data into addressable fast write storage and/or as part of writing data into allocation units associated with erase blocks. Erasure codes may be used across storage devices, as well as within erase blocks or allocation units, or within and across Flash memory devices on a single storage device, to provide redundancy against single or multiple storage device failures or to protect against internal corruptions of Flash memory pages resulting from Flash memory operations or from degradation of Flash memory cells. Mirroring and erasure coding at various levels may be used to recover from multiple types of failures that occur separately or in combination.

[0080] The embodiments depicted with reference to FIGS. 2A-G illustrate a storage cluster that stores user data, such as user data originating from one or more user or client systems or other sources external to the storage cluster. The storage cluster distributes user data across storage nodes housed within a chassis, or across multiple chassis, using erasure coding and redundant copies of metadata. Erasure coding refers to a method of data protection or reconstruction in which data is stored across a set of different locations, such as disks, storage nodes or geographic locations. Flash memory is one type of solid-state memory that may be integrated with the embodiments, although the embodiments may be extended to other types of solid-state memory or other storage medium, including non-solid state memory. Control of storage locations and workloads are distributed across the storage locations in a clustered peer-to-peer system. Tasks such as mediating communications between the various storage nodes, detecting when a storage node has become unavailable, and balancing I/Os (inputs and outputs) across the various storage nodes, are all handled on a distributed basis. Data is laid out or distributed across multiple storage nodes in data fragments or stripes that support data recovery in some embodiments. Ownership of data can be reassigned within a cluster, independent of input and output patterns. This architecture described in more detail below allows a storage node in the cluster to fail, with the system remaining operational, since the data can be reconstructed from other storage nodes and thus remain available for input and output operations. In various embodiments, a storage node may be referred to as a cluster node, a blade, or a server.

[0081] The storage cluster may be contained within a chassis, i.e., an enclosure housing one or more storage nodes. A mechanism to provide power to each storage node, such as a power distribution bus, and a communication mechanism, such as a communication bus that enables communication between the storage nodes are included within the chassis. The storage cluster can run as an independent system in one location according to some embodiments. In one embodiment, a chassis contains at least two instances of both the power distribution and the communication bus which may be enabled or disabled independently. The internal communication bus may be an Ethernet bus, however, other technologies such as PCIe, InfiniBand, and others, are equally suitable. The chassis provides a port for an external communication bus for enabling communication between multiple chassis, directly or through a switch, and with client systems. The external communication may use a technology such as Ethernet, InfiniBand, Fibre Channel, etc. In some embodiments, the external communication bus uses

different communication bus technologies for inter-chassis and client communication. If a switch is deployed within or between chassis, the switch may act as a translation between multiple protocols or technologies. When multiple chassis are connected to define a storage cluster, the storage cluster may be accessed by a client using either proprietary interfaces or standard interfaces such as network file system ('NFS'), common internet file system ('CIFS'), small computer system interface ('SCSI') or hypertext transfer protocol ('HTTP'). Translation from the client protocol may occur at the switch, chassis external communication bus or within each storage node. In some embodiments, multiple chassis may be coupled or connected to each other through an aggregator switch. A portion and/or all of the coupled or connected chassis may be designated as a storage cluster. As discussed above, each chassis can have multiple blades, each blade has a media access control ('MAC') address, but the storage cluster is presented to an external network as having a single cluster IP address and a single MAC address in some embodiments.

[0082] Each storage node may be one or more storage servers and each storage server is connected to one or more non-volatile solid state memory units, which may be referred to as storage units or storage devices. One embodiment includes a single storage server in each storage node and between one to eight non-volatile solid state memory units, however this one example is not meant to be limiting. The storage server may include a processor, DRAM and interfaces for the internal communication bus and power distribution for each of the power buses. Inside the storage node, the interfaces and storage unit share a communication bus, e.g., PCI Express, in some embodiments. The non-volatile solid state memory units may directly access the internal communication bus interface through a storage node communication bus, or request the storage node to access the bus interface. The non-volatile solid state memory unit contains an embedded CPU, solid state storage controller, and a quantity of solid state mass storage, e.g., between 2-32 terabytes ('TB') in some embodiments. An embedded volatile storage medium, such as DRAM, and an energy reserve apparatus are included in the non-volatile solid state memory unit. In some embodiments, the energy reserve apparatus is a capacitor, super-capacitor, or battery that enables transferring a subset of DRAM contents to a stable storage medium in the case of power loss. In some embodiments, the non-volatile solid state memory unit is constructed with a storage class memory, such as phase change or magnetoresistive random access memory ('MRAM') that substitutes for DRAM and enables a reduced power hold-up apparatus.

[0083] One of many features of the storage nodes and non-volatile solid state storage is the ability to proactively rebuild data in a storage cluster. The storage nodes and non-volatile solid state storage can determine when a storage node or non-volatile solid state storage in the storage cluster is unreachable, independent of whether there is an attempt to read data involving that storage node or non-volatile solid state storage. The storage nodes and non-volatile solid state storage then cooperate to recover and rebuild the data in at least partially new locations. This constitutes a proactive rebuild, in that the system rebuilds data without waiting until the data is needed for a read access initiated from a client

system employing the storage cluster. These and further details of the storage memory and operation thereof are discussed below.

[0084] FIG. 2A is a perspective view of a storage cluster **161**, with multiple storage nodes **150** and internal solid-state memory coupled to each storage node to provide network attached storage or storage area network, in accordance with some embodiments. A network attached storage, storage area network, or a storage cluster, or other storage memory, could include one or more storage clusters **161**, each having one or more storage nodes **150**, in a flexible and reconfigurable arrangement of both the physical components and the amount of storage memory provided thereby. The storage cluster **161** is designed to fit in a rack, and one or more racks can be set up and populated as desired for the storage memory. The storage cluster **161** has a chassis **138** having multiple slots **142**. It should be appreciated that chassis **138** may be referred to as a housing, enclosure, or rack unit. In one embodiment, the chassis **138** has fourteen slots **142**, although other numbers of slots are readily devised. For example, some embodiments have four slots, eight slots, sixteen slots, thirty-two slots, or other suitable number of slots. Each slot **142** can accommodate one storage node **150** in some embodiments. Chassis **138** includes flaps **148** that can be utilized to mount the chassis **138** on a rack. Fans **144** provide air circulation for cooling of the storage nodes **150** and components thereof, although other cooling components could be used, or an embodiment could be devised without cooling components. A switch fabric **146** couples storage nodes **150** within chassis **138** together and to a network for communication to the memory. In an embodiment depicted in herein, the slots **142** to the left of the switch fabric **146** and fans **144** are shown occupied by storage nodes **150**, while the slots **142** to the right of the switch fabric **146** and fans **144** are empty and available for insertion of storage node **150** for illustrative purposes. This configuration is one example, and one or more storage nodes **150** could occupy the slots **142** in various further arrangements. The storage node arrangements need not be sequential or adjacent in some embodiments. Storage nodes **150** are hot pluggable, meaning that a storage node **150** can be inserted into a slot **142** in the chassis **138**, or removed from a slot **142**, without stopping or powering down the system. Upon insertion or removal of storage node **150** from slot **142**, the system automatically reconfigures in order to recognize and adapt to the change. Reconfiguration, in some embodiments, includes restoring redundancy and/or rebalancing data or load.

[0085] Each storage node **150** can have multiple components. In the embodiment shown here, the storage node **150** includes a printed circuit board **159** populated by a CPU **156**, i.e., processor, a memory **154** coupled to the CPU **156**, and a non-volatile solid state storage **152** coupled to the CPU **156**, although other mountings and/or components could be used in further embodiments. The memory **154** has instructions which are executed by the CPU **156** and/or data operated on by the CPU **156**. As further explained below, the non-volatile solid state storage **152** includes flash or, in further embodiments, other types of solid-state memory.

[0086] Referring to FIG. 2A, storage cluster **161** is scalable, meaning that storage capacity with non-uniform storage sizes is readily added, as described above. One or more storage nodes **150** can be plugged into or removed from each chassis and the storage cluster self-configures in some

embodiments. Plug-in storage nodes **150**, whether installed in a chassis as delivered or later added, can have different sizes. For example, in one embodiment a storage node **150** can have any multiple of 4 TB, e.g., 8 TB, 12 TB, 16 TB, 32 TB, etc. In further embodiments, a storage node **150** could have any multiple of other storage amounts or capacities. Storage capacity of each storage node **150** is broadcast, and influences decisions of how to stripe the data. For maximum storage efficiency, an embodiment can self-configure as wide as possible in the stripe, subject to a predetermined requirement of continued operation with loss of up to one, or up to two, non-volatile solid state storage **152** units or storage nodes **150** within the chassis.

[0087] FIG. 2B is a block diagram showing a communications interconnect **173** and power distribution bus **172** coupling multiple storage nodes **150**. Referring back to FIG. 2A, the communications interconnect **173** can be included in or implemented with the switch fabric **146** in some embodiments. Where multiple storage clusters **161** occupy a rack, the communications interconnect **173** can be included in or implemented with a top of rack switch, in some embodiments. As illustrated in FIG. 2B, storage cluster **161** is enclosed within a single chassis **138**. External port **176** is coupled to storage nodes **150** through communications interconnect **173**, while external port **174** is coupled directly to a storage node. External power port **178** is coupled to power distribution bus **172**. Storage nodes **150** may include varying amounts and differing capacities of non-volatile solid state storage **152** as described with reference to FIG. 2A. In addition, one or more storage nodes **150** may be a compute only storage node as illustrated in FIG. 2B. Authorities **168** are implemented on the non-volatile solid state storage **152**, for example as lists or other data structures stored in memory. In some embodiments the authorities are stored within the non-volatile solid state storage **152** and supported by software executing on a controller or other processor of the non-volatile solid state storage **152**. In a further embodiment, authorities **168** are implemented on the storage nodes **150**, for example as lists or other data structures stored in the memory **154** and supported by software executing on the CPU **156** of the storage node **150**. Authorities **168** control how and where data is stored in the non-volatile solid state storage **152** in some embodiments. This control assists in determining which type of erasure coding scheme is applied to the data, and which storage nodes **150** have which portions of the data. Each authority **168** may be assigned to a non-volatile solid state storage **152**. Each authority may control a range of inode numbers, segment numbers, or other data identifiers which are assigned to data by a file system, by the storage nodes **150**, or by the non-volatile solid state storage **152**, in various embodiments.

[0088] Every piece of data, and every piece of metadata, has redundancy in the system in some embodiments. In addition, every piece of data and every piece of metadata has an owner, which may be referred to as an authority. If that authority is unreachable, for example through failure of a storage node, there is a plan of succession for how to find that data or that metadata. In various embodiments, there are redundant copies of authorities **168**. Authorities **168** have a relationship to storage nodes **150** and non-volatile solid state storage **152** in some embodiments. Each authority **168**, covering a range of data segment numbers or other identifiers of the data, may be assigned to a specific non-volatile solid state storage **152**. In some embodiments the authorities

168 for all of such ranges are distributed over the non-volatile solid state storage **152** of a storage cluster. Each storage node **150** has a network port that provides access to the non-volatile solid state storage(s) **152** of that storage node **150**. Data can be stored in a segment, which is associated with a segment number and that segment number is an indirection for a configuration of a RAID (redundant array of independent disks) stripe in some embodiments. The assignment and use of the authorities **168** thus establishes an indirection to data. Indirection may be referred to as the ability to reference data indirectly, in this case via an authority **168**, in accordance with some embodiments. A segment identifies a set of non-volatile solid state storage **152** and a local identifier into the set of non-volatile solid state storage **152** that may contain data. In some embodiments, the local identifier is an offset into the device and may be reused sequentially by multiple segments. In other embodiments the local identifier is unique for a specific segment and never reused. The offsets in the non-volatile solid state storage **152** are applied to locating data for writing to or reading from the non-volatile solid state storage **152** (in the form of a RAID stripe). Data is striped across multiple units of non-volatile solid state storage **152**, which may include or be different from the non-volatile solid state storage **152** having the authority **168** for a particular data segment.

[0089] If there is a change in where a particular segment of data is located, e.g., during a data move or a data reconstruction, the authority **168** for that data segment should be consulted, at that non-volatile solid state storage **152** or storage node **150** having that authority **168**. In order to locate a particular piece of data, embodiments calculate a hash value for a data segment or apply an inode number or a data segment number. The output of this operation points to a non-volatile solid state storage **152** having the authority **168** for that particular piece of data. In some embodiments there are two stages to this operation. The first stage maps an entity identifier (ID), e.g., a segment number, inode number, or directory number to an authority identifier. This mapping may include a calculation such as a hash or a bit mask. The second stage is mapping the authority identifier to a particular non-volatile solid state storage **152**, which may be done through an explicit mapping. The operation is repeatable, so that when the calculation is performed, the result of the calculation repeatably and reliably points to a particular non-volatile solid state storage **152** having that authority **168**. The operation may include the set of reachable storage nodes as input. If the set of reachable non-volatile solid state storage units changes the optimal set changes. In some embodiments, the persisted value is the current assignment (which is always true) and the calculated value is the target assignment the cluster will attempt to reconfigure towards. This calculation may be used to determine the optimal non-volatile solid state storage **152** for an authority in the presence of a set of non-volatile solid state storage **152** that are reachable and constitute the same cluster. The calculation also determines an ordered set of peer non-volatile solid state storage **152** that will also record the authority to non-volatile solid state storage mapping so that the authority may be determined even if the assigned non-volatile solid state storage is unreachable. A duplicate or substitute authority **168** may be consulted if a specific authority **168** is unavailable in some embodiments.

[0090] With reference to FIGS. 2A and 2B, two of the many tasks of the CPU 156 on a storage node 150 are to break up write data, and reassemble read data. When the system has determined that data is to be written, the authority 168 for that data is located as above. When the segment ID for data is already determined the request to write is forwarded to the non-volatile solid state storage 152 currently determined to be the host of the authority 168 determined from the segment. The host CPU 156 of the storage node 150, on which the non-volatile solid state storage 152 and corresponding authority 168 reside, then breaks up or shards the data and transmits the data out to various non-volatile solid state storage 152. The transmitted data is written as a data stripe in accordance with an erasure coding scheme. In some embodiments, data is requested to be pulled, and in other embodiments, data is pushed. In reverse, when data is read, the authority 168 for the segment ID containing the data is located as described above. The host CPU 156 of the storage node 150 on which the non-volatile solid state storage 152 and corresponding authority 168 reside requests the data from the non-volatile solid state storage and corresponding storage nodes pointed to by the authority. In some embodiments the data is read from flash storage as a data stripe. The host CPU 156 of storage node 150 then reassembles the read data, correcting any errors (if present) according to the appropriate erasure coding scheme, and forwards the reassembled data to the network. In further embodiments, some or all of these tasks can be handled in the non-volatile solid state storage 152. In some embodiments, the segment host requests the data be sent to storage node 150 by requesting pages from storage and then sending the data to the storage node making the original request.

[0091] In embodiments, authorities 168 operate to determine how operations will proceed against particular logical elements. Each of the logical elements may be operated on through a particular authority across a plurality of storage controllers of a storage system. The authorities 168 may communicate with the plurality of storage controllers so that the plurality of storage controllers collectively perform operations against those particular logical elements.

[0092] In embodiments, logical elements could be, for example, files, directories, object buckets, individual objects, delineated parts of files or objects, other forms of key-value pair databases, or tables. In embodiments, performing an operation can involve, for example, ensuring consistency, structural integrity, and/or recoverability with other operations against the same logical element, reading metadata and data associated with that logical element, determining what data should be written durably into the storage system to persist any changes for the operation, or where metadata and data can be determined to be stored across modular storage devices attached to a plurality of the storage controllers in the storage system.

[0093] In some embodiments the operations are token based transactions to efficiently communicate within a distributed system. Each transaction may be accompanied by or associated with a token, which gives permission to execute the transaction. The authorities 168 are able to maintain a pre-transaction state of the system until completion of the operation in some embodiments. The token based communication may be accomplished without a global lock across the system, and also enables restart of an operation in case of a disruption or other failure.

[0094] In some systems, for example in UNIX-style file systems, data is handled with an index node or inode, which specifies a data structure that represents an object in a file system. The object could be a file or a directory, for example. Metadata may accompany the object, as attributes such as permission data and a creation timestamp, among other attributes. A segment number could be assigned to all or a portion of such an object in a file system. In other systems, data segments are handled with a segment number assigned elsewhere. For purposes of discussion, the unit of distribution is an entity, and an entity can be a file, a directory or a segment. That is, entities are units of data or metadata stored by a storage system. Entities are grouped into sets called authorities. Each authority has an authority owner, which is a storage node that has the exclusive right to update the entities in the authority. In other words, a storage node contains the authority, and that the authority, in turn, contains entities.

[0095] A segment is a logical container of data in accordance with some embodiments. A segment is an address space between medium address space and physical flash locations, i.e., the data segment number, are in this address space. Segments may also contain meta-data, which enable data redundancy to be restored (rewritten to different flash locations or devices) without the involvement of higher level software. In one embodiment, an internal format of a segment contains client data and medium mappings to determine the position of that data. Each data segment is protected, e.g., from memory and other failures, by breaking the segment into a number of data and parity shards, where applicable. The data and parity shards are distributed, i.e., striped, across non-volatile solid state storage 152 coupled to the host CPUs 156 (See FIGS. 2E and 2G) in accordance with an erasure coding scheme. Usage of the term segments refers to the container and its place in the address space of segments in some embodiments. Usage of the term stripe refers to the same set of shards as a segment and includes how the shards are distributed along with redundancy or parity information in accordance with some embodiments.

[0096] A series of address-space transformations takes place across an entire storage system. At the top are the directory entries (file names) which link to an inode. Inodes point into medium address space, where data is logically stored. Medium addresses may be mapped through a series of indirect mediums to spread the load of large files, or implement data services like deduplication or snapshots. Medium addresses may be mapped through a series of indirect mediums to spread the load of large files, or implement data services like deduplication or snapshots. Segment addresses are then translated into physical flash locations. Physical flash locations have an address range bounded by the amount of flash in the system in accordance with some embodiments. Medium addresses and segment addresses are logical containers, and in some embodiments use a 128 bit or larger identifier so as to be practically infinite, with a likelihood of reuse calculated as longer than the expected life of the system. Addresses from logical containers are allocated in a hierarchical fashion in some embodiments. Initially, each non-volatile solid state storage 152 unit may be assigned a range of address space. Within this assigned range, the non-volatile solid state storage 152 is able to allocate addresses without synchronization with other non-volatile solid state storage 152.

[0097] Data and metadata is stored by a set of underlying storage layouts that are optimized for varying workload patterns and storage devices. These layouts incorporate multiple redundancy schemes, compression formats and index algorithms. Some of these layouts store information about authorities and authority masters, while others store file metadata and file data. The redundancy schemes include error correction codes that tolerate corrupted bits within a single storage device (such as a NAND flash chip), erasure codes that tolerate the failure of multiple storage nodes, and replication schemes that tolerate data center or regional failures. In some embodiments, low density parity check ('LDPC') code is used within a single storage unit. Reed-Solomon encoding is used within a storage cluster, and mirroring is used within a storage grid in some embodiments. Metadata may be stored using an ordered log structured index (such as a Log Structured Merge Tree), and large data may not be stored in a log structured layout.

[0098] In order to maintain consistency across multiple copies of an entity, the storage nodes agree implicitly on two things through calculations: (1) the authority that contains the entity, and (2) the storage node that contains the authority. The assignment of entities to authorities can be done by pseudo randomly assigning entities to authorities, by splitting entities into ranges based upon an externally produced key, or by placing a single entity into each authority. Examples of pseudorandom schemes are linear hashing and the Replication Under Scalable Hashing ('RUSH') family of hashes, including Controlled Replication Under Scalable Hashing ('CRUSH'). In some embodiments, pseudo-random assignment is utilized only for assigning authorities to nodes because the set of nodes can change. The set of authorities cannot change so any subjective function may be applied in these embodiments. Some placement schemes automatically place authorities on storage nodes, while other placement schemes rely on an explicit mapping of authorities to storage nodes. In some embodiments, a pseudorandom scheme is utilized to map from each authority to a set of candidate authority owners. A pseudorandom data distribution function related to CRUSH may assign authorities to storage nodes and create a list of where the authorities are assigned. Each storage node has a copy of the pseudorandom data distribution function, and can arrive at the same calculation for distributing, and later finding or locating an authority. Each of the pseudorandom schemes requires the reachable set of storage nodes as input in some embodiments in order to conclude the same target nodes. Once an entity has been placed in an authority, the entity may be stored on physical devices so that no expected failure will lead to unexpected data loss. In some embodiments, rebalancing algorithms attempt to store the copies of all entities within an authority in the same layout and on the same set of machines.

[0099] Examples of expected failures include device failures, stolen machines, datacenter fires, and regional disasters, such as nuclear or geological events. Different failures lead to different levels of acceptable data loss. In some embodiments, a stolen storage node impacts neither the security nor the reliability of the system, while depending on system configuration, a regional event could lead to no loss of data, a few seconds or minutes of lost updates, or even complete data loss.

[0100] In the embodiments, the placement of data for storage redundancy is independent of the placement of

authorities for data consistency. In some embodiments, storage nodes that contain authorities do not contain any persistent storage. Instead, the storage nodes are connected to non-volatile solid state storage units that do not contain authorities. The communications interconnect between storage nodes and non-volatile solid state storage units consists of multiple communication technologies and has non-uniform performance and fault tolerance characteristics. In some embodiments, as mentioned above, non-volatile solid state storage units are connected to storage nodes via PCI express, storage nodes are connected together within a single chassis using Ethernet backplane, and chassis are connected together to form a storage cluster. Storage clusters are connected to clients using Ethernet or fiber channel in some embodiments. If multiple storage clusters are configured into a storage grid, the multiple storage clusters are connected using the Internet or other long-distance networking links, such as a "metro scale" link or private link that does not traverse the internet.

[0101] Authority owners have the exclusive right to modify entities, to migrate entities from one non-volatile solid state storage unit to another non-volatile solid state storage unit, and to add and remove copies of entities. This allows for maintaining the redundancy of the underlying data. When an authority owner fails, is going to be decommissioned, or is overloaded, the authority is transferred to a new storage node. Transient failures make it non-trivial to ensure that all non-faulty machines agree upon the new authority location. The ambiguity that arises due to transient failures can be achieved automatically by a consensus protocol such as Paxos, hot-warm failover schemes, via manual intervention by a remote system administrator, or by a local hardware administrator (such as by physically removing the failed machine from the cluster, or pressing a button on the failed machine). In some embodiments, a consensus protocol is used, and failover is automatic. If too many failures or replication events occur in too short a time period, the system goes into a self-preservation mode and halts replication and data movement activities until an administrator intervenes in accordance with some embodiments.

[0102] As authorities are transferred between storage nodes and authority owners update entities in their authorities, the system transfers messages between the storage nodes and non-volatile solid state storage units. With regard to persistent messages, messages that have different purposes are of different types. Depending on the type of the message, the system maintains different ordering and durability guarantees. As the persistent messages are being processed, the messages are temporarily stored in multiple durable and non-durable storage hardware technologies. In some embodiments, messages are stored in RAM, NVRAM and on NAND flash devices, and a variety of protocols are used in order to make efficient use of each storage medium. Latency-sensitive client requests may be persisted in replicated NVRAM, and then later NAND, while background rebalancing operations are persisted directly to NAND.

[0103] Persistent messages are persistently stored prior to being transmitted. This allows the system to continue to serve client requests despite failures and component replacement. Although many hardware components contain unique identifiers that are visible to system administrators, manufacturer, hardware supply chain and ongoing monitoring quality control infrastructure, applications running on top of

the infrastructure address virtualize addresses. These virtualized addresses do not change over the lifetime of the storage system, regardless of component failures and replacements. This allows each component of the storage system to be replaced over time without reconfiguration or disruptions of client request processing, i.e., the system supports non-disruptive upgrades.

[0104] In some embodiments, the virtualized addresses are stored with sufficient redundancy. A continuous monitoring system correlates hardware and software status and the hardware identifiers. This allows detection and prediction of failures due to faulty components and manufacturing details. The monitoring system also enables the proactive transfer of authorities and entities away from impacted devices before failure occurs by removing the component from the critical path in some embodiments.

[0105] FIG. 2C is a multiple level block diagram, showing contents of a storage node **150** and contents of a non-volatile solid state storage **152** of the storage node **150**. Data is communicated to and from the storage node **150** by a network interface controller ('NIC') **202** in some embodiments. Each storage node **150** has a CPU **156**, and one or more non-volatile solid state storage **152**, as discussed above. Moving down one level in FIG. 2C, each non-volatile solid state storage **152** has a relatively fast non-volatile solid state memory, such as nonvolatile random access memory ('NVRAM') **204**, and flash memory **206**. In some embodiments, NVRAM **204** may be a component that does not require program/erase cycles (DRAM, MRAM, PCM), and can be a memory that can support being written vastly more often than the memory is read from. Moving down another level in FIG. 2C, the NVRAM **204** is implemented in one embodiment as high speed volatile memory, such as dynamic random access memory (DRAM) **216**, backed up by energy reserve **218**. Energy reserve **218** provides sufficient electrical power to keep the DRAM **216** powered long enough for contents to be transferred to the flash memory **206** in the event of power failure. In some embodiments, energy reserve **218** is a capacitor, super-capacitor, battery, or other device, that supplies a suitable supply of energy sufficient to enable the transfer of the contents of DRAM **216** to a stable storage medium in the case of power loss. The flash memory **206** is implemented as multiple flash dies **222**, which may be referred to as packages of flash dies **222** or an array of flash dies **222**. It should be appreciated that the flash dies **222** could be packaged in any number of ways, with a single die per package, multiple dies per package (i.e., multichip packages), in hybrid packages, as bare dies on a printed circuit board or other substrate, as encapsulated dies, etc. In the embodiment shown, the non-volatile solid state storage **152** has a controller **212** or other processor, and an input output (I/O) port **210** coupled to the controller **212**. I/O port **210** is coupled to the CPU **156** and/or the network interface controller **202** of the flash storage node **150**. Flash input output (I/O) port **220** is coupled to the flash dies **222**, and a direct memory access unit (DMA) **214** is coupled to the controller **212**, the DRAM **216** and the flash dies **222**. In the embodiment shown, the I/O port **210**, controller **212**, DMA unit **214** and flash I/O port **220** are implemented on a programmable logic device ('PLD') **208**, e.g., an FPGA. In this embodiment, each flash die **222** has pages, organized as sixteen kB (kilobyte) pages **224**, and a register **226** through which data can be written to or read from the flash die **222**.

In further embodiments, other types of solid-state memory are used in place of, or in addition to flash memory illustrated within flash die **222**.

[0106] Storage clusters **161**, in various embodiments as disclosed herein, can be contrasted with storage arrays in general. The storage nodes **150** are part of a collection that creates the storage cluster **161**. Each storage node **150** owns a slice of data and computing required to provide the data. Multiple storage nodes **150** cooperate to store and retrieve the data. Storage memory or storage devices, as used in storage arrays in general, are less involved with processing and manipulating the data. Storage memory or storage devices in a storage array receive commands to read, write, or erase data. The storage memory or storage devices in a storage array are not aware of a larger system in which they are embedded, or what the data means. Storage memory or storage devices in storage arrays can include various types of storage memory, such as RAM, solid state drives, hard disk drives, etc. The non-volatile solid state storage **152** units described herein have multiple interfaces active simultaneously and serving multiple purposes. In some embodiments, some of the functionality of a storage node **150** is shifted into a storage unit **152**, transforming the storage unit **152** into a combination of storage unit **152** and storage node **150**. Placing computing (relative to storage data) into the storage unit **152** places this computing closer to the data itself. The various system embodiments have a hierarchy of storage node layers with different capabilities. By contrast, in a storage array, a controller owns and knows everything about all of the data that the controller manages in a shelf or storage devices. In a storage cluster **161**, as described herein, multiple controllers in multiple non-volatile solid state storage **152** units and/or storage nodes **150** cooperate in various ways (e.g., for erasure coding, data sharding, metadata communication and redundancy, storage capacity expansion or contraction, data recovery, and so on).

[0107] FIG. 2D shows a storage server environment, which uses embodiments of the storage nodes **150** and storage **152** units of FIGS. 2A-C. In this version, each non-volatile solid state storage **152** unit has a processor such as controller **212** (see FIG. 2C), an FPGA, flash memory **206**, and NVRAM **204** (which is super-capacitor backed DRAM **216**, see FIGS. 2B and 2C) on a PCIe (peripheral component interconnect express) board in a chassis **138** (see FIG. 2A). The non-volatile solid state storage **152** unit may be implemented as a single board containing storage, and may be the largest tolerable failure domain inside the chassis. In some embodiments, up to two non-volatile solid state storage **152** units may fail and the device will continue with no data loss.

[0108] The physical storage is divided into named regions based on application usage in some embodiments. The NVRAM **204** is a contiguous block of reserved memory in the non-volatile solid state storage **152** DRAM **216**, and is backed by NAND flash. NVRAM **204** is logically divided into multiple memory regions written for two as spool (e.g., spool region). Space within the NVRAM **204** spools is managed by each authority **168** independently. Each device provides an amount of storage space to each authority **168**. That authority **168** further manages lifetimes and allocations within that space. Examples of a spool include distributed transactions or notions. When the primary power to a non-volatile solid state storage **152** unit fails, onboard super-capacitors provide a short duration of power hold up. During

this holdup interval, the contents of the NVRAM 204 are flushed to flash memory 206. On the next power-on, the contents of the NVRAM 204 are recovered from the flash memory 206.

[0109] As for the storage unit controller, the responsibility of the logical “controller” is distributed across each of the blades containing authorities 168. This distribution of logical control is shown in FIG. 2D as a host controller 242, mid-tier controller 244 and storage unit controller(s) 246. Management of the control plane and the storage plane are treated independently, although parts may be physically co-located on the same blade. Each authority 168 effectively serves as an independent controller. Each authority 168 provides its own data and metadata structures, its own background workers, and maintains its own lifecycle.

[0110] FIG. 2E is a blade 252 hardware block diagram, showing a control plane 254, compute and storage planes 256, 258, and authorities 168 interacting with underlying physical resources, using embodiments of the storage nodes 150 and storage units 152 of FIGS. 2A-C in the storage server environment of FIG. 2D. The control plane 254 is partitioned into a number of authorities 168 which can use the compute resources in the compute plane 256 to run on any of the blades 252. The storage plane 258 is partitioned into a set of devices, each of which provides access to flash 206 and NVRAM 204 resources. In one embodiment, the compute plane 256 may perform the operations of a storage array controller, as described herein, on one or more devices of the storage plane 258 (e.g., a storage array).

[0111] In the compute and storage planes 256, 258 of FIG. 2E, the authorities 168 interact with the underlying physical resources (i.e., devices). From the point of view of an authority 168, its resources are striped over all of the physical devices. From the point of view of a device, it provides resources to all authorities 168, irrespective of where the authorities happen to run. Each authority 168 has allocated or has been allocated one or more partitions 260 of storage memory in the storage units 152, e.g., partitions 260 in flash memory 206 and NVRAM 204. Each authority 168 uses those allocated partitions 260 that belong to it, for writing or reading user data. Authorities can be associated with differing amounts of physical storage of the system. For example, one authority 168 could have a larger number of partitions 260 or larger sized partitions 260 in one or more storage units 152 than one or more other authorities 168.

[0112] FIG. 2F depicts elasticity software layers in blades 252 of a storage cluster, in accordance with some embodiments. In the elasticity structure, elasticity software is symmetric, i.e., each blade’s compute module 270 runs the three identical layers of processes depicted in FIG. 2F. Storage managers 274 execute read and write requests from other blades 252 for data and metadata stored in local storage unit 152 NVRAM 204 and flash 206. Authorities 168 fulfill client requests by issuing the necessary reads and writes to the blades 252 on whose storage units 152 the corresponding data or metadata resides. Endpoints 272 parse client connection requests received from switch fabric 146 supervisory software, relay the client connection requests to the authorities 168 responsible for fulfillment, and relay the authorities’ 168 responses to clients. The symmetric three-layer structure enables the storage system’s high degree of concurrency. Elasticity scales out efficiently and reliably in these embodiments. In addition, elasticity implements a unique scale-out technique that balances work evenly across

all resources regardless of client access pattern, and maximizes concurrency by eliminating much of the need for inter-blade coordination that typically occurs with conventional distributed locking.

[0113] Still referring to FIG. 2F, authorities 168 running in the compute modules 270 of a blade 252 perform the internal operations required to fulfill client requests. One feature of elasticity is that authorities 168 are stateless, i.e., they cache active data and metadata in their own blades’ 252 DRAMs for fast access, but the authorities store every update in their NVRAM 204 partitions on three separate blades 252 until the update has been written to flash 206. All the storage system writes to NVRAM 204 are in triplicate to partitions on three separate blades 252 in some embodiments. With triple-mirrored NVRAM 204 and persistent storage protected by parity and Reed-Solomon RAID checksums, the storage system can survive concurrent failure of two blades 252 with no loss of data, metadata, or access to either.

[0114] Because authorities 168 are stateless, they can migrate between blades 252. Each authority 168 has a unique identifier. NVRAM 204 and flash 206 partitions are associated with authorities’ 168 identifiers, not with the blades 252 on which they are running in some. Thus, when an authority 168 migrates, the authority 168 continues to manage the same storage partitions from its new location. When a new blade 252 is installed in an embodiment of the storage cluster, the system automatically rebalances load by: partitioning the new blade’s 252 storage for use by the system’s authorities 168, migrating selected authorities 168 to the new blade 252, starting endpoints 272 on the new blade 252 and including them in the switch fabric’s 146 client connection distribution algorithm.

[0115] From their new locations, migrated authorities 168 persist the contents of their NVRAM 204 partitions on flash 206, process read and write requests from other authorities 168, and fulfill the client requests that endpoints 272 direct to them. Similarly, if a blade 252 fails or is removed, the system redistributes its authorities 168 among the system’s remaining blades 252. The redistributed authorities 168 continue to perform their original functions from their new locations.

[0116] FIG. 2G depicts authorities 168 and storage resources in blades 252 of a storage cluster, in accordance with some embodiments. Each authority 168 is exclusively responsible for a partition of the flash 206 and NVRAM 204 on each blade 252. The authority 168 manages the content and integrity of its partitions independently of other authorities 168. Authorities 168 compress incoming data and preserve it temporarily in their NVRAM 204 partitions, and then consolidate, RAID-protect, and persist the data in segments of the storage in their flash 206 partitions. As the authorities 168 write data to flash 206, storage managers 274 perform the necessary flash translation to optimize write performance and maximize media longevity. In the background, authorities 168 “garbage collect,” or reclaim space occupied by data that clients have made obsolete by overwriting the data. It should be appreciated that since authorities’ 168 partitions are disjoint, there is no need for distributed locking to execute client and writes or to perform background functions.

[0117] The embodiments described herein may utilize various software, communication and/or networking protocols. In addition, the configuration of the hardware and/or software may be adjusted to accommodate various proto-

cols. For example, the embodiments may utilize Active Directory, which is a database based system that provides authentication, directory, policy, and other services in a WINDOWSTTM environment. In these embodiments, LDAP (Lightweight Directory Access Protocol) is one example application protocol for querying and modifying items in directory service providers such as Active Directory. In some embodiments, a network lock manager ('NLM') is utilized as a facility that works in cooperation with the Network File System ('NFS') to provide a System V style of advisory file and record locking over a network. The Server Message Block ('SMB') protocol, one version of which is also known as Common Internet File System ('CIFS'), may be integrated with the storage systems discussed herein. SMP operates as an application-layer network protocol typically used for providing shared access to files, printers, and serial ports and miscellaneous communications between nodes on a network. SMB also provides an authenticated inter-process communication mechanism. AMAZONTM S3 (Simple Storage Service) is a web service offered by Amazon Web Services, and the systems described herein may interface with Amazon S3 through web services interfaces (REST (representational state transfer), SOAP (simple object access protocol), and BitTorrent). A RESTful API (application programming interface) breaks down a transaction to create a series of small modules. Each module addresses a particular underlying part of the transaction. The control or permissions provided with these embodiments, especially for object data, may include utilization of an access control list ('ACL'). The ACL is a list of permissions attached to an object and the ACL specifies which users or system processes are granted access to objects, as well as what operations are allowed on given objects. The systems may utilize Internet Protocol version 6 ('IPv6'), as well as IPv4, for the communications protocol that provides an identification and location system for computers on networks and routes traffic across the Internet. The routing of packets between networked systems may include Equal-cost multi-path routing ('ECMP'), which is a routing strategy where next-hop packet forwarding to a single destination can occur over multiple "best paths" which tie for top place in routing metric calculations. Multi-path routing can be used in conjunction with most routing protocols, because it is a per-hop decision limited to a single router. The software may support Multi-tenancy, which is an architecture in which a single instance of a software application serves multiple customers. Each customer may be referred to as a tenant. Tenants may be given the ability to customize some parts of the application, but may not customize the application's code, in some embodiments. The embodiments may maintain audit logs. An audit log is a document that records an event in a computing system. In addition to documenting what resources were accessed, audit log entries typically include destination and source addresses, a timestamp, and user login information for compliance with various regulations. The embodiments may support various key management policies, such as encryption key rotation. In addition, the system may support dynamic root passwords or some variation dynamically changing passwords.

[0118] FIG. 3A sets forth a diagram of a storage system 306 that is coupled for data communications with a cloud services provider 302 in accordance with some embodiments of the present disclosure. Although depicted in less detail, the storage system 306 depicted in FIG. 3A may be

similar to the storage systems described above with reference to FIGS. 1A-1D and FIGS. 2A-2G. In some embodiments, the storage system 306 depicted in FIG. 3A may be embodied as a storage system that includes imbalanced active/active controllers, as a storage system that includes balanced active/active controllers, as a storage system that includes active/active controllers where less than all of each controller's resources are utilized such that each controller has reserve resources that may be used to support failover, as a storage system that includes fully active/active controllers, as a storage system that includes dataset-segregated controllers, as a storage system that includes dual-layer architectures with front-end controllers and back-end integrated storage controllers, as a storage system that includes scale-out clusters of dual-controller arrays, as well as combinations of such embodiments.

[0119] In the example depicted in FIG. 3A, the storage system 306 is coupled to the cloud services provider 302 via a data communications link 304. The data communications link 304 may be embodied as a dedicated data communications link, as a data communications pathway that is provided through the use of one or data communications networks such as a wide area network ('WAN') or LAN, or as some other mechanism capable of transporting digital information between the storage system 306 and the cloud services provider 302. Such a data communications link 304 may be fully wired, fully wireless, or some aggregation of wired and wireless data communications pathways. In such an example, digital information may be exchanged between the storage system 306 and the cloud services provider 302 via the data communications link 304 using one or more data communications protocols. For example, digital information may be exchanged between the storage system 306 and the cloud services provider 302 via the data communications link 304 using the handheld device transfer protocol ('HDTP'), hypertext transfer protocol ('HTTP'), internet protocol ('IP'), real-time transfer protocol ('RTP'), transmission control protocol ('TCP'), user datagram protocol ('UDP'), wireless application protocol ('WAP'), or other protocol.

[0120] The cloud services provider 302 depicted in FIG. 3A may be embodied, for example, as a system and computing environment that provides a vast array of services to users of the cloud services provider 302 through the sharing of computing resources via the data communications link 304. The cloud services provider 302 may provide on-demand access to a shared pool of configurable computing resources such as computer networks, servers, storage, applications and services, and so on. The shared pool of configurable resources may be rapidly provisioned and released to a user of the cloud services provider 302 with minimal management effort. Generally, the user of the cloud services provider 302 is unaware of the exact computing resources utilized by the cloud services provider 302 to provide the services. Although in many cases such a cloud services provider 302 may be accessible via the Internet, readers of skill in the art will recognize that any system that abstracts the use of shared resources to provide services to a user through any data communications link may be considered a cloud services provider 302.

[0121] In the example depicted in FIG. 3A, the cloud services provider 302 may be configured to provide a variety of services to the storage system 306 and users of the storage system 306 through the implementation of various service

models. For example, the cloud services provider **302** may be configured to provide services through the implementation of an infrastructure as a service ('IaaS') service model, through the implementation of a platform as a service ('PaaS') service model, through the implementation of a software as a service ('SaaS') service model, through the implementation of an authentication as a service ('AaaS') service model, through the implementation of a storage as a service model where the cloud services provider **302** offers access to its storage infrastructure for use by the storage system **306** and users of the storage system **306**, and so on. Readers will appreciate that the cloud services provider **302** may be configured to provide additional services to the storage system **306** and users of the storage system **306** through the implementation of additional service models, as the service models described above are included only for explanatory purposes and in no way represent a limitation of the services that may be offered by the cloud services provider **302** or a limitation as to the service models that may be implemented by the cloud services provider **302**.

[0122] In the example depicted in FIG. 3A, the cloud services provider **302** may be embodied, for example, as a private cloud, as a public cloud, or as a combination of a private cloud and public cloud. In an embodiment in which the cloud services provider **302** is embodied as a private cloud, the cloud services provider **302** may be dedicated to providing services to a single organization rather than providing services to multiple organizations. In an embodiment where the cloud services provider **302** is embodied as a public cloud, the cloud services provider **302** may provide services to multiple organizations. In still alternative embodiments, the cloud services provider **302** may be embodied as a mix of a private and public cloud services with a hybrid cloud deployment.

[0123] Although not explicitly depicted in FIG. 3A, readers will appreciate that a vast amount of additional hardware components and additional software components may be necessary to facilitate the delivery of cloud services to the storage system **306** and users of the storage system **306**. For example, the storage system **306** may be coupled to (or even include) a cloud storage gateway. Such a cloud storage gateway may be embodied, for example, as hardware-based or software-based appliance that is located on premise with the storage system **306**. Such a cloud storage gateway may operate as a bridge between local applications that are executing on the storage system **306** and remote, cloud-based storage that is utilized by the storage system **306**. Through the use of a cloud storage gateway, organizations may move primary iSCSI or NAS to the cloud services provider **302**, thereby enabling the organization to save space on their on-premises storage systems. Such a cloud storage gateway may be configured to emulate a disk array, a block-based device, a file server, or other storage system that can translate the SCSI commands, file server commands, or other appropriate command into REST-space protocols that facilitate communications with the cloud services provider **302**.

[0124] In order to enable the storage system **306** and users of the storage system **306** to make use of the services provided by the cloud services provider **302**, a cloud migration process may take place during which data, applications, or other elements from an organization's local systems (or even from another cloud environment) are moved to the cloud services provider **302**. In order to successfully migrate

data, applications, or other elements to the cloud services provider's **302** environment, middleware such as a cloud migration tool may be utilized to bridge gaps between the cloud services provider's **302** environment and an organization's environment. Such cloud migration tools may also be configured to address potentially high network costs and long transfer times associated with migrating large volumes of data to the cloud services provider **302**, as well as addressing security concerns associated with sensitive data to the cloud services provider **302** over data communications networks. In order to further enable the storage system **306** and users of the storage system **306** to make use of the services provided by the cloud services provider **302**, a cloud orchestrator may also be used to arrange and coordinate automated tasks in pursuit of creating a consolidated process or workflow. Such a cloud orchestrator may perform tasks such as configuring various components, whether those components are cloud components or on-premises components, as well as managing the interconnections between such components. The cloud orchestrator can simplify the inter-component communication and connections to ensure that links are correctly configured and maintained.

[0125] In the example depicted in FIG. 3A, and as described briefly above, the cloud services provider **302** may be configured to provide services to the storage system **306** and users of the storage system **306** through the usage of a SaaS service model, eliminating the need to install and run the application on local computers, which may simplify maintenance and support of the application. Such applications may take many forms in accordance with various embodiments of the present disclosure. For example, the cloud services provider **302** may be configured to provide access to data analytics applications to the storage system **306** and users of the storage system **306**. Such data analytics applications may be configured, for example, to receive vast amounts of telemetry data phoned home by the storage system **306**. Such telemetry data may describe various operating characteristics of the storage system **306** and may be analyzed for a vast array of purposes including, for example, to determine the health of the storage system **306**, to identify workloads that are executing on the storage system **306**, to predict when the storage system **306** will run out of various resources, to recommend configuration changes, hardware or software upgrades, workflow migrations, or other actions that may improve the operation of the storage system **306**.

[0126] The cloud services provider **302** may also be configured to provide access to virtualized computing environments to the storage system **306** and users of the storage system **306**. Such virtualized computing environments may be embodied, for example, as a virtual machine or other virtualized computer hardware platforms, virtual storage devices, virtualized computer network resources, and so on. Examples of such virtualized environments can include virtual machines that are created to emulate an actual computer, virtualized desktop environments that separate a logical desktop from a physical machine, virtualized file systems that allow uniform access to different types of concrete file systems, and many others.

[0127] Although the example depicted in FIG. 3A illustrates the storage system **306** being coupled for data communications with the cloud services provider **302**, in other embodiments the storage system **306** may be part of a hybrid cloud deployment in which private cloud elements (e.g.,

private cloud services, on-premises infrastructure, and so on) and public cloud elements (e.g., public cloud services, infrastructure, and so on that may be provided by one or more cloud services providers) are combined to form a single solution, with orchestration among the various platforms. Such a hybrid cloud deployment may leverage hybrid cloud management software such as, for example, Azure™ Arc from Microsoft™, that centralize the management of the hybrid cloud deployment to any infrastructure and enable the deployment of services anywhere. In such an example, the hybrid cloud management software may be configured to create, update, and delete resources (both physical and virtual) that form the hybrid cloud deployment, to allocate compute and storage to specific workloads, to monitor workloads and resources for performance, policy compliance, updates and patches, security status, or to perform a variety of other tasks.

[0128] Readers will appreciate that by pairing the storage systems described herein with one or more cloud services providers, various offerings may be enabled. For example, disaster recovery as a service ('DRaaS') may be provided where cloud resources are utilized to protect applications and data from disruption caused by disaster, including in embodiments where the storage systems may serve as the primary data store. In such embodiments, a total system backup may be taken that allows for business continuity in the event of system failure. In such embodiments, cloud data backup techniques (by themselves or as part of a larger DRaaS solution) may also be integrated into an overall solution that includes the storage systems and cloud services providers described herein.

[0129] The storage systems described herein, as well as the cloud services providers, may be utilized to provide a wide array of security features. For example, the storage systems may encrypt data at rest (and data may be sent to and from the storage systems encrypted) and may make use of Key Management-as-a-Service ('KMaaS') to manage encryption keys, keys for locking and unlocking storage devices, and so on. Likewise, cloud data security gateways or similar mechanisms may be utilized to ensure that data stored within the storage systems does not improperly end up being stored in the cloud as part of a cloud data backup operation. Furthermore, microsegmentation or identity-based-segmentation may be utilized in a data center that includes the storage systems or within the cloud services provider, to create secure zones in data centers and cloud deployments that enables the isolation of workloads from one another.

[0130] For further explanation, FIG. 3B sets forth a diagram of a storage system 306 in accordance with some embodiments of the present disclosure. Although depicted in less detail, the storage system 306 depicted in FIG. 3B may be similar to the storage systems described above with reference to FIGS. 1A-1D and FIGS. 2A-2G as the storage system may include many of the components described above.

[0131] The storage system 306 depicted in FIG. 3B may include a vast amount of storage resources 308, which may be embodied in many forms. For example, the storage resources 308 can include nano-RAM or another form of nonvolatile random access memory that utilizes carbon nanotubes deposited on a substrate, 3D crosspoint non-volatile memory, flash memory including single-level cell ('SLC') NAND flash, multi-level cell ('MLC') NAND flash,

triple-level cell ('TLC') NAND flash, quad-level cell ('QLC') NAND flash, or others. Likewise, the storage resources 308 may include non-volatile magnetoresistive random-access memory ('MRAM'), including spin transfer torque ('STT') MRAM. The example storage resources 308 may alternatively include non-volatile phase-change memory ('PCM'), quantum memory that allows for the storage and retrieval of photonic quantum information, resistive random-access memory ('ReRAM'), storage class memory ('SCM'), or other form of storage resources, including any combination of resources described herein. Readers will appreciate that other forms of computer memories and storage devices may be utilized by the storage systems described above, including DRAM, SRAM, EEPROM, universal memory, and many others. The storage resources 308 depicted in FIG. 3A may be embodied in a variety of form factors, including but not limited to, dual in-line memory modules ('DIMMs'), non-volatile dual in-line memory modules ('NVDIMMs'), M.2, U.2, and others.

[0132] The storage resources 308 depicted in FIG. 3B may include various forms of SCM. SCM may effectively treat fast, non-volatile memory (e.g., NAND flash) as an extension of DRAM such that an entire dataset may be treated as an in-memory dataset that resides entirely in DRAM. SCM may include non-volatile media such as, for example, NAND flash. Such NAND flash may be accessed utilizing NVMe that can use the PCIe bus as its transport, providing for relatively low access latencies compared to older protocols. In fact, the network protocols used for SSDs in all-flash arrays can include NVMe using Ethernet (ROCE, NVMe TCP), Fibre Channel (NVMe FC), InfiniBand (iWARP), and others that make it possible to treat fast, non-volatile memory as an extension of DRAM. In view of the fact that DRAM is often byte-addressable and fast, non-volatile memory such as NAND flash is block-addressable, a controller software/hardware stack may be needed to convert the block data to the bytes that are stored in the media. Examples of media and software that may be used as SCM can include, for example, 3D XPoint, Intel Memory Drive Technology, Samsung's Z-SSD, and others.

[0133] The storage resources 308 depicted in FIG. 3B may also include racetrack memory (also referred to as domain-wall memory). Such racetrack memory may be embodied as a form of non-volatile, solid-state memory that relies on the intrinsic strength and orientation of the magnetic field created by an electron as it spins in addition to its electronic charge, in solid-state devices. Through the use of spin-coherent electric current to move magnetic domains along a nanoscopic permalloy wire, the domains may pass by magnetic read/write heads positioned near the wire as current is passed through the wire, which alter the domains to record patterns of bits. In order to create a racetrack memory device, many such wires and read/write elements may be packaged together.

[0134] The example storage system 306 depicted in FIG. 3B may implement a variety of storage architectures. For example, storage systems in accordance with some embodiments of the present disclosure may utilize block storage where data is stored in blocks, and each block essentially acts as an individual hard drive. Storage systems in accordance with some embodiments of the present disclosure may utilize object storage, where data is managed as objects. Each object may include the data itself, a variable amount of metadata, and a globally unique identifier, where object

storage can be implemented at multiple levels (e.g., device level, system level, interface level). Storage systems in accordance with some embodiments of the present disclosure utilize file storage in which data is stored in a hierarchical structure. Such data may be saved in files and folders, and presented to both the system storing it and the system retrieving it in the same format.

[0135] The example storage system 306 depicted in FIG. 3B may be embodied as a storage system in which additional storage resources can be added through the use of a scale-up model, additional storage resources can be added through the use of a scale-out model, or through some combination thereof. In a scale-up model, additional storage may be added by adding additional storage devices. In a scale-out model, however, additional storage nodes may be added to a cluster of storage nodes, where such storage nodes can include additional processing resources, additional networking resources, and so on.

[0136] The example storage system 306 depicted in FIG. 3B may leverage the storage resources described above in a variety of different ways. For example, some portion of the storage resources may be utilized to serve as a write cache, storage resources within the storage system may be utilized as a read cache, or tiering may be achieved within the storage systems by placing data within the storage system in accordance with one or more tiering policies.

[0137] The storage system 306 depicted in FIG. 3B also includes communications resources 310 that may be useful in facilitating data communications between components within the storage system 306, as well as data communications between the storage system 306 and computing devices that are outside of the storage system 306, including embodiments where those resources are separated by a relatively vast expanse. The communications resources 310 may be configured to utilize a variety of different protocols and data communication fabrics to facilitate data communications between components within the storage systems as well as computing devices that are outside of the storage system. For example, the communications resources 310 can include fibre channel ('FC') technologies such as FC fabrics and FC protocols that can transport SCSI commands over FC network, FC over ethernet ('FCoE') technologies through which FC frames are encapsulated and transmitted over Ethernet networks, InfiniBand ('IB') technologies in which a switched fabric topology is utilized to facilitate transmissions between channel adapters, NVM Express ('NVMe') technologies and NVMe over fabrics ('NVMeoF') technologies through which non-volatile storage media attached via a PCI express ('PCIe') bus may be accessed, and others. In fact, the storage systems described above may, directly or indirectly, make use of neutrino communication technologies and devices through which information (including binary information) is transmitted using a beam of neutrinos.

[0138] The communications resources 310 can also include mechanisms for accessing storage resources 308 within the storage system 306 utilizing serial attached SCSI ('SAS'), serial ATA ('SATA') bus interfaces for connecting storage resources 308 within the storage system 306 to host bus adapters within the storage system 306, internet small computer systems interface ('iSCSI') technologies to provide block-level access to storage resources 308 within the storage system 306, and other communications resources that may be useful in facilitating data communications between components within the storage system 306, as well

as data communications between the storage system 306 and computing devices that are outside of the storage system 306.

[0139] The storage system 306 depicted in FIG. 3B also includes processing resources 312 that may be useful in useful in executing computer program instructions and performing other computational tasks within the storage system 306. The processing resources 312 may include one or more ASICs that are customized for some particular purpose as well as one or more CPUs. The processing resources 312 may also include one or more DSPs, one or more FPGAs, one or more systems on a chip ('SoCs'), or other form of processing resources 312. The storage system 306 may utilize the storage resources 312 to perform a variety of tasks including, but not limited to, supporting the execution of software resources 314 that will be described in greater detail below.

[0140] The storage system 306 depicted in FIG. 3B also includes software resources 314 that, when executed by processing resources 312 within the storage system 306, may perform a vast array of tasks. The software resources 314 may include, for example, one or more modules of computer program instructions that when executed by processing resources 312 within the storage system 306 are useful in carrying out various data protection techniques. Such data protection techniques may be carried out, for example, by system software executing on computer hardware within the storage system, by a cloud services provider, or in other ways. Such data protection techniques can include data archiving, data backup, data replication, data snapshotting, data and database cloning, and other data protection techniques.

[0141] The software resources 314 may also include software that is useful in implementing software-defined storage ('SDS'). In such an example, the software resources 314 may include one or more modules of computer program instructions that, when executed, are useful in policy-based provisioning and management of data storage that is independent of the underlying hardware. Such software resources 314 may be useful in implementing storage virtualization to separate the storage hardware from the software that manages the storage hardware.

[0142] The software resources 314 may also include software that is useful in facilitating and optimizing I/O operations that are directed to the storage system 306. For example, the software resources 314 may include software modules that perform various data reduction techniques such as, for example, data compression, data deduplication, and others. The software resources 314 may include software modules that intelligently group together I/O operations to facilitate better usage of the underlying storage resource 308, software modules that perform data migration operations to migrate from within a storage system, as well as software modules that perform other functions. Such software resources 314 may be embodied as one or more software containers or in many other ways.

[0143] For further explanation, FIG. 3C sets forth an example of a cloud-based storage system 318 in accordance with some embodiments of the present disclosure. In the example depicted in FIG. 3C, the cloud-based storage system 318 is created entirely in a cloud computing environment 316 such as, for example, Amazon Web Services ('AWS')TM, Microsoft AzureTM, Google Cloud PlatformTM, IBM CloudTM, Oracle CloudTM, and others. The cloud-based

storage system **318** may be used to provide services similar to the services that may be provided by the storage systems described above.

[0144] The cloud-based storage system **318** depicted in FIG. 3C includes two cloud computing instances **320**, **322** that each are used to support the execution of a storage controller application **324**, **326**. The cloud computing instances **320**, **322** may be embodied, for example, as instances of cloud computing resources (e.g., virtual machines) that may be provided by the cloud computing environment **316** to support the execution of software applications such as the storage controller application **324**, **326**. For example, each of the cloud computing instances **320**, **322** may execute on an Azure VM, where each Azure VM may include high speed temporary storage that may be leveraged as a cache (e.g., as a read cache). In one embodiment, the cloud computing instances **320**, **322** may be embodied as Amazon Elastic Compute Cloud ('EC2') instances. In such an example, an Amazon Machine Image ('AMI') that includes the storage controller application **324**, **326** may be booted to create and configure a virtual machine that may execute the storage controller application **324**, **326**.

[0145] In the example method depicted in FIG. 3C, the storage controller application **324**, **326** may be embodied as a module of computer program instructions that, when executed, carries out various storage tasks. For example, the storage controller application **324**, **326** may be embodied as a module of computer program instructions that, when executed, carries out the same tasks as the controllers **110A**, **110B** in FIG. 1A described above such as writing data to the cloud-based storage system **318**, erasing data from the cloud-based storage system **318**, retrieving data from the cloud-based storage system **318**, monitoring and reporting of disk utilization and performance, performing redundancy operations, such as RAID or RAID-like data redundancy operations, compressing data, encrypting data, deduplicating data, and so forth. Readers will appreciate that because there are two cloud computing instances **320**, **322** that each include the storage controller application **324**, **326**, in some embodiments one cloud computing instance **320** may operate as the primary controller as described above while the other cloud computing instance **322** may operate as the secondary controller as described above. Readers will appreciate that the storage controller application **324**, **326** depicted in FIG. 3C may include identical source code that is executed within different cloud computing instances **320**, **322** such as distinct EC2 instances.

[0146] Readers will appreciate that other embodiments that do not include a primary and secondary controller are within the scope of the present disclosure. For example, each cloud computing instance **320**, **322** may operate as a primary controller for some portion of the address space supported by the cloud-based storage system **318**, each cloud computing instance **320**, **322** may operate as a primary controller where the servicing of I/O operations directed to the cloud-based storage system **318** are divided in some other way, and so on. In fact, in other embodiments where costs savings may be prioritized over performance demands, only a single cloud computing instance may exist that contains the storage controller application.

[0147] The cloud-based storage system **318** depicted in FIG. 3C includes cloud computing instances **340a**, **340b**, **340n** with local storage **330**, **334**, **338**. The cloud computing instances **340a**, **340b**, **340n** may be embodied, for example,

as instances of cloud computing resources that may be provided by the cloud computing environment **316** to support the execution of software applications. The cloud computing instances **340a**, **340b**, **340n** of FIG. 3C may differ from the cloud computing instances **320**, **322** described above as the cloud computing instances **340a**, **340b**, **340n** of FIG. 3C have local storage **330**, **334**, **338** resources whereas the cloud computing instances **320**, **322** that support the execution of the storage controller application **324**, **326** need not have local storage resources. The cloud computing instances **340a**, **340b**, **340n** with local storage **330**, **334**, **338** may be embodied, for example, as EC2 M5 instances that include one or more SSDs, as EC2 R5 instances that include one or more SSDs, as EC2 I3 instances that include one or more SSDs, and so on. In some embodiments, the local storage **330**, **334**, **338** must be embodied as solid-state storage (e.g., SSDs) rather than storage that makes use of hard disk drives.

[0148] In the example depicted in FIG. 3C, each of the cloud computing instances **340a**, **340b**, **340n** with local storage **330**, **334**, **338** can include a software daemon **328**, **332**, **336** that, when executed by a cloud computing instance **340a**, **340b**, **340n** can present itself to the storage controller applications **324**, **326** as if the cloud computing instance **340a**, **340b**, **340n** were a physical storage device (e.g., one or more SSDs). In such an example, the software daemon **328**, **332**, **336** may include computer program instructions similar to those that would normally be contained on a storage device such that the storage controller applications **324**, **326** can send and receive the same commands that a storage controller would send to storage devices. In such a way, the storage controller applications **324**, **326** may include code that is identical to (or substantially identical to) the code that would be executed by the controllers in the storage systems described above. In these and similar embodiments, communications between the storage controller applications **324**, **326** and the cloud computing instances **340a**, **340b**, **340n** with local storage **330**, **334**, **338** may utilize iSCSI, NVMe over TCP, messaging, a custom protocol, or in some other mechanism.

[0149] In the example depicted in FIG. 3C, each of the cloud computing instances **340a**, **340b**, **340n** with local storage **330**, **334**, **338** may also be coupled to block storage **342**, **344**, **346** that is offered by the cloud computing environment **316** such as, for example, as Amazon Elastic Block Store ('EBS') volumes. In such an example, the block storage **342**, **344**, **346** that is offered by the cloud computing environment **316** may be utilized in a manner that is similar to how the NVRAM devices described above are utilized, as the software daemon **328**, **332**, **336** (or some other module) that is executing within a particular cloud computing instance **340a**, **340b**, **340n** may, upon receiving a request to write data, initiate a write of the data to its attached EBS volume as well as a write of the data to its local storage **330**, **334**, **338** resources. In some alternative embodiments, data may only be written to the local storage **330**, **334**, **338** resources within a particular cloud computing instance **340a**, **340b**, **340n**. In an alternative embodiment, rather than using the block storage **342**, **344**, **346** that is offered by the cloud computing environment **316** as NVRAM, actual RAM on each of the cloud computing instances **340a**, **340b**, **340n** with local storage **330**, **334**, **338** may be used as NVRAM, thereby decreasing network utilization costs that would be associated with using an EBS volume as the NVRAM. In yet

another embodiment, high performance block storage resources such as one or more Azure Ultra Disks may be utilized as the NVRAM.

[0150] The storage controller applications 324, 326 may be used to perform various tasks such as deduplicating the data contained in the request, compressing the data contained in the request, determining where to write the data contained in the request, and so on, before ultimately sending a request to write a deduplicated, encrypted, or otherwise possibly updated version of the data to one or more of the cloud computing instances 340a, 340b, 340n with local storage 330, 334, 338. Either cloud computing instance 320, 322, in some embodiments, may receive a request to read data from the cloud-based storage system 318 and may ultimately send a request to read data to one or more of the cloud computing instances 340a, 340b, 340n with local storage 330, 334, 338.

[0151] When a request to write data is received by a particular cloud computing instance 340a, 340b, 340n with local storage 330, 334, 338, the software daemon 328, 332, 336 may be configured to not only write the data to its own local storage 330, 334, 338 resources and any appropriate block storage 342, 344, 346 resources, but the software daemon 328, 332, 336 may also be configured to write the data to cloud-based object storage 348 that is attached to the particular cloud computing instance 340a, 340b, 340n. The cloud-based object storage 348 that is attached to the particular cloud computing instance 340a, 340b, 340n may be embodied, for example, as Amazon Simple Storage Service ('S3'). In other embodiments, the cloud computing instances 320, 322 that each include the storage controller application 324, 326 may initiate the storage of the data in the local storage 330, 334, 338 of the cloud computing instances 340a, 340b, 340n and the cloud-based object storage 348. In other embodiments, rather than using both the cloud computing instances 340a, 340b, 340n with local storage 330, 334, 338 (also referred to herein as 'virtual drives') and the cloud-based object storage 348 to store data, a persistent storage layer may be implemented in other ways. For example, one or more Azure Ultra disks may be used to persistently store data (e.g., after the data has been written to the NVRAM layer).

[0152] While the local storage 330, 334, 338 resources and the block storage 342, 344, 346 resources that are utilized by the cloud computing instances 340a, 340b, 340n may support block-level access, the cloud-based object storage 348 that is attached to the particular cloud computing instance 340a, 340b, 340n supports only object-based access. The software daemon 328, 332, 336 may therefore be configured to take blocks of data, package those blocks into objects, and write the objects to the cloud-based object storage 348 that is attached to the particular cloud computing instance 340a, 340b, 340n.

[0153] Consider an example in which data is written to the local storage 330, 334, 338 resources and the block storage 342, 344, 346 resources that are utilized by the cloud computing instances 340a, 340b, 340n in 1 MB blocks. In such an example, assume that a user of the cloud-based storage system 318 issues a request to write data that, after being compressed and deduplicated by the storage controller application 324, 326 results in the need to write 5 MB of data. In such an example, writing the data to the local storage 330, 334, 338 resources and the block storage 342, 344, 346 resources that are utilized by the cloud computing instances

340a, 340b, 340n is relatively straightforward as 5 blocks that are 1 MB in size are written to the local storage 330, 334, 338 resources and the block storage 342, 344, 346 resources that are utilized by the cloud computing instances 340a, 340b, 340n. In such an example, the software daemon 328, 332, 336 may also be configured to create five objects containing distinct 1 MB chunks of the data. As such, in some embodiments, each object that is written to the cloud-based object storage 348 may be identical (or nearly identical) in size. Readers will appreciate that in such an example, metadata that is associated with the data itself may be included in each object (e.g., the first 1 MB of the object is data and the remaining portion is metadata associated with the data). Readers will appreciate that the cloud-based object storage 348 may be incorporated into the cloud-based storage system 318 to increase the durability of the cloud-based storage system 318.

[0154] In some embodiments, all data that is stored by the cloud-based storage system 318 may be stored in both: 1) the cloud-based object storage 348, and 2) at least one of the local storage 330, 334, 338 resources or block storage 342, 344, 346 resources that are utilized by the cloud computing instances 340a, 340b, 340n. In such embodiments, the local storage 330, 334, 338 resources and block storage 342, 344, 346 resources that are utilized by the cloud computing instances 340a, 340b, 340n may effectively operate as cache that generally includes all data that is also stored in S3, such that all reads of data may be serviced by the cloud computing instances 340a, 340b, 340n without requiring the cloud computing instances 340a, 340b, 340n to access the cloud-based object storage 348. Readers will appreciate that in other embodiments, however, all data that is stored by the cloud-based storage system 318 may be stored in the cloud-based object storage 348, but less than all data that is stored by the cloud-based storage system 318 may be stored in at least one of the local storage 330, 334, 338 resources or block storage 342, 344, 346 resources that are utilized by the cloud computing instances 340a, 340b, 340n. In such an example, various policies may be utilized to determine which subset of the data that is stored by the cloud-based storage system 318 should reside in both: 1) the cloud-based object storage 348, and 2) at least one of the local storage 330, 334, 338 resources or block storage 342, 344, 346 resources that are utilized by the cloud computing instances 340a, 340b, 340n.

[0155] One or more modules of computer program instructions that are executing within the cloud-based storage system 318 (e.g., a monitoring module that is executing on its own EC2 instance) may be designed to handle the failure of one or more of the cloud computing instances 340a, 340b, 340n with local storage 330, 334, 338. In such an example, the monitoring module may handle the failure of one or more of the cloud computing instances 340a, 340b, 340n with local storage 330, 334, 338 by creating one or more new cloud computing instances with local storage, retrieving data that was stored on the failed cloud computing instances 340a, 340b, 340n from the cloud-based object storage 348, and storing the data retrieved from the cloud-based object storage 348 in local storage on the newly created cloud computing instances. Readers will appreciate that many variants of this process may be implemented.

[0156] Readers will appreciate that various performance aspects of the cloud-based storage system 318 may be monitored (e.g., by a monitoring module that is executing in

an EC2 instance) such that the cloud-based storage system **318** can be scaled-up or scaled-out as needed. For example, if the cloud computing instances **320**, **322** that are used to support the execution of a storage controller application **324**, **326** are undersized and not sufficiently servicing the I/O requests that are issued by users of the cloud-based storage system **318**, a monitoring module may create a new, more powerful cloud computing instance (e.g., a cloud computing instance of a type that includes more processing power, more memory, etc. . . .) that includes the storage controller application such that the new, more powerful cloud computing instance can begin operating as the primary controller. Likewise, if the monitoring module determines that the cloud computing instances **320**, **322** that are used to support the execution of a storage controller application **324**, **326** are oversized and that cost savings could be gained by switching to a smaller, less powerful cloud computing instance, the monitoring module may create a new, less powerful (and less expensive) cloud computing instance that includes the storage controller application such that the new, less powerful cloud computing instance can begin operating as the primary controller.

[0157] The storage systems described above may carry out intelligent data backup techniques through which data stored in the storage system may be copied and stored in a distinct location to avoid data loss in the event of equipment failure or some other form of catastrophe. For example, the storage systems described above may be configured to examine each backup to avoid restoring the storage system to an undesirable state. Consider an example in which malware infects the storage system. In such an example, the storage system may include software resources **314** that can scan each backup to identify backups that were captured before the malware infected the storage system and those backups that were captured after the malware infected the storage system. In such an example, the storage system may restore itself from a backup that does not include the malware—or at least not restore the portions of a backup that contained the malware. In such an example, the storage system may include software resources **314** that can scan each backup to identify the presences of malware (or a virus, or some other undesirable), for example, by identifying write operations that were serviced by the storage system and originated from a network subnet that is suspected to have delivered the malware, by identifying write operations that were serviced by the storage system and originated from a user that is suspected to have delivered the malware, by identifying write operations that were serviced by the storage system and examining the content of the write operation against fingerprints of the malware, and in many other ways.

[0158] Readers will further appreciate that the backups (often in the form of one or more snapshots) may also be utilized to perform rapid recovery of the storage system. Consider an example in which the storage system is infected with ransomware that locks users out of the storage system. In such an example, software resources **314** within the storage system may be configured to detect the presence of ransomware and may be further configured to restore the storage system to a point-in-time, using the retained backups, prior to the point-in-time at which the ransomware infected the storage system. In such an example, the presence of ransomware may be explicitly detected through the use of software tools utilized by the system, through the use of a key (e.g., a USB drive) that is inserted into the storage

system, or in a similar way. Likewise, the presence of ransomware may be inferred in response to system activity meeting a predetermined fingerprint such as, for example, no reads or writes coming into the system for a predetermined period of time.

[0159] Readers will appreciate that the various components described above may be grouped into one or more optimized computing packages as converged infrastructures. Such converged infrastructures may include pools of computers, storage and networking resources that can be shared by multiple applications and managed in a collective manner using policy-driven processes. Such converged infrastructures may be implemented with a converged infrastructure reference architecture, with standalone appliances, with a software driven hyper-converged approach (e.g., hyper-converged infrastructures), or in other ways.

[0160] Readers will appreciate that the storage systems described in this disclosure may be useful for supporting various types of software applications. In fact, the storage systems may be ‘application aware’ in the sense that the storage systems may obtain, maintain, or otherwise have access to information describing connected applications (e.g., applications that utilize the storage systems) to optimize the operation of the storage system based on intelligence about the applications and their utilization patterns. For example, the storage system may optimize data layouts, optimize caching behaviors, optimize ‘QoS’ levels, or perform some other optimization that is designed to improve the storage performance that is experienced by the application.

[0161] As an example of one type of application that may be supported by the storage systems describe herein, the storage system **306** may be useful in supporting artificial intelligence (‘AI’) applications, database applications, XOps projects (e.g., DevOps projects, DataOps projects, MLOps projects, ModelOps projects, PlatformOps projects), electronic design automation tools, event-driven software applications, high performance computing applications, simulation applications, high-speed data capture and analysis applications, machine learning applications, media production applications, media serving applications, picture archiving and communication systems (‘PACS’) applications, software development applications, virtual reality applications, augmented reality applications, and many other types of applications by providing storage resources to such applications.

[0162] In view of the fact that the storage systems include compute resources, storage resources, and a wide variety of other resources, the storage systems may be well suited to support applications that are resource intensive such as, for example, AI applications. AI applications may be deployed in a variety of fields, including: predictive maintenance in manufacturing and related fields, healthcare applications such as patient data & risk analytics, retail and marketing deployments (e.g., search advertising, social media advertising), supply chains solutions, fintech solutions such as business analytics & reporting tools, operational deployments such as real-time analytics tools, application performance management tools, IT infrastructure management tools, and many others.

[0163] Such AI applications may enable devices to perceive their environment and take actions that maximize their chance of success at some goal. Examples of such AI

applications can include IBM Watson™, Microsoft Oxford™, Google DeepMind™, Baidu Minwa™, and others.

[0164] The storage systems described above may also be well suited to support other types of applications that are resource intensive such as, for example, machine learning applications. Machine learning applications may perform various types of data analysis to automate analytical model building. Using algorithms that iteratively learn from data, machine learning applications can enable computers to learn without being explicitly programmed. One particular area of machine learning is referred to as reinforcement learning, which involves taking suitable actions to maximize reward in a particular situation.

[0165] In addition to the resources already described, the storage systems described above may also include graphics processing units ('GPUs'), occasionally referred to as visual processing unit ('VDUs'). Such GPUs may be embodied as specialized electronic circuits that rapidly manipulate and alter memory to accelerate the creation of images in a frame buffer intended for output to a display device. Such GPUs may be included within any of the computing devices that are part of the storage systems described above, including as one of many individually scalable components of a storage system, where other examples of individually scalable components of such storage system can include storage components, memory components, compute components (e.g., CPUs, FPGAs, ASICs), networking components, software components, and others. In addition to GPUs, the storage systems described above may also include neural network processors ('NNPs') for use in various aspects of neural network processing. Such NNPs may be used in place of (or in addition to) GPUs and may also be independently scalable.

[0166] As described above, the storage systems described herein may be configured to support artificial intelligence applications, machine learning applications, big data analytics applications, and many other types of applications. The rapid growth in these sort of applications is being driven by three technologies: deep learning (DL), GPU processors, and Big Data. Deep learning is a computing model that makes use of massively parallel neural networks inspired by the human brain. Instead of experts handcrafting software, a deep learning model writes its own software by learning from lots of examples. Such GPUs may include thousands of cores that are well-suited to run algorithms that loosely represent the parallel nature of the human brain.

[0167] Advances in deep neural networks, including the development of multi-layer neural networks, have ignited a new wave of algorithms and tools for data scientists to tap into their data with artificial intelligence (AI). With improved algorithms, larger data sets, and various frameworks (including open-source software libraries for machine learning across a range of tasks), data scientists are tackling new use cases like autonomous driving vehicles, natural language processing and understanding, computer vision, machine reasoning, strong AI, and many others. Applications of such techniques may include: machine and vehicular object detection, identification and avoidance; visual recognition, classification and tagging; algorithmic financial trading strategy performance management; simultaneous localization and mapping; predictive maintenance of high-value machinery; prevention against cyber security threats, expertise automation; image recognition and classification; ques-

tion answering; robotics; text analytics (extraction, classification) and text generation and translation; and many others. Applications of AI techniques has materialized in a wide array of products include, for example, Amazon Echo's speech recognition technology that allows users to talk to their machines, Google Translate™ which allows for machine-based language translation, Spotify's Discover Weekly that provides recommendations on new songs and artists that a user may like based on the user's usage and traffic analysis, Quill's text generation offering that takes structured data and turns it into narrative stories, Chatbots that provide real-time, contextually specific answers to questions in a dialog format, and many others.

[0168] Data is the heart of modern AI and deep learning algorithms. Before training can begin, one problem that must be addressed revolves around collecting the labeled data that is crucial for training an accurate AI model. A full scale AI deployment may be required to continuously collect, clean, transform, label, and store large amounts of data. Adding additional high quality data points directly translates to more accurate models and better insights. Data samples may undergo a series of processing steps including, but not limited to: 1) ingesting the data from an external source into the training system and storing the data in raw form, 2) cleaning and transforming the data in a format convenient for training, including linking data samples to the appropriate label, 3) exploring parameters and models, quickly testing with a smaller dataset, and iterating to converge on the most promising models to push into the production cluster, 4) executing training phases to select random batches of input data, including both new and older samples, and feeding those into production GPU servers for computation to update model parameters, and 5) evaluating including using a holdback portion of the data not used in training in order to evaluate model accuracy on the holdout data. This lifecycle may apply for any type of parallelized machine learning, not just neural networks or deep learning. For example, standard machine learning frameworks may rely on CPUs instead of GPUs but the data ingest and training workflows may be the same. Readers will appreciate that a single shared storage data hub creates a coordination point throughout the lifecycle without the need for extra data copies among the ingest, preprocessing, and training stages. Rarely is the ingested data used for only one purpose, and shared storage gives the flexibility to train multiple different models or apply traditional analytics to the data.

[0169] Readers will appreciate that each stage in the AI data pipeline may have varying requirements from the data hub (e.g., the storage system or collection of storage systems). Scale-out storage systems must deliver uncompromising performance for all manner of access types and patterns—from small, metadata-heavy to large files, from random to sequential access patterns, and from low to high concurrency. The storage systems described above may serve as an ideal AI data hub as the systems may service unstructured workloads. In the first stage, data is ideally ingested and stored on to the same data hub that following stages will use, in order to avoid excess data copying. The next two steps can be done on a standard compute server that optionally includes a GPU, and then in the fourth and last stage, full training production jobs are run on powerful GPU-accelerated servers. Often, there is a production pipeline alongside an experimental pipeline operating on the same dataset. Further, the GPU-accelerated servers can be

used independently for different models or joined together to train on one larger model, even spanning multiple systems for distributed training. If the shared storage tier is slow, then data must be copied to local storage for each phase, resulting in wasted time staging data onto different servers. The ideal data hub for the AI training pipeline delivers performance similar to data stored locally on the server node while also having the simplicity and performance to enable all pipeline stages to operate concurrently.

[0170] In order for the storage systems described above to serve as a data hub or as part of an AI deployment, in some embodiments the storage systems may be configured to provide DMA between storage devices that are included in the storage systems and one or more GPUs that are used in an AI or big data analytics pipeline. The one or more GPUs may be coupled to the storage system, for example, via NVMe-over-Fabrics (‘NVMe-oF’) such that bottlenecks such as the host CPU can be bypassed and the storage system (or one of the components contained therein) can directly access GPU memory. In such an example, the storage systems may leverage API hooks to the GPUs to transfer data directly to the GPUs. For example, the GPUs may be embodied as Nvidia™ GPUs and the storage systems may support GPUDirect Storage (‘GDS’) software, or have similar proprietary software, that enables the storage system to transfer data to the GPUs via RDMA or similar mechanism.

[0171] Although the preceding paragraphs discuss deep learning applications, readers will appreciate that the storage systems described herein may also be part of a distributed deep learning (‘DDL’) platform to support the execution of DDL algorithms. The storage systems described above may also be paired with other technologies such as TensorFlow, an open-source software library for dataflow programming across a range of tasks that may be used for machine learning applications such as neural networks, to facilitate the development of such machine learning models, applications, and so on.

[0172] The storage systems described above may also be used in a neuromorphic computing environment. Neuromorphic computing is a form of computing that mimics brain cells. To support neuromorphic computing, an architecture of interconnected “neurons” replace traditional computing models with low-powered signals that go directly between neurons for more efficient computation. Neuromorphic computing may make use of very-large-scale integration (VLSI) systems containing electronic analog circuits to mimic neuro-biological architectures present in the nervous system, as well as analog, digital, mixed-mode analog/digital VLSI, and software systems that implement models of neural systems for perception, motor control, or multisensory integration.

[0173] Readers will appreciate that the storage systems described above may be configured to support the storage or use of (among other types of data) blockchains and derivative items such as, for example, open source blockchains and related tools that are part of the IBM™ Hyperledger project, permissioned blockchains in which a certain number of trusted parties are allowed to access the block chain, blockchain products that enable developers to build their own distributed ledger projects, and others. Blockchains and the storage systems described herein may be leveraged to support on-chain storage of data as well as off-chain storage of data.

[0174] Off-chain storage of data can be implemented in a variety of ways and can occur when the data itself is not stored within the blockchain. For example, in one embodiment, a hash function may be utilized and the data itself may be fed into the hash function to generate a hash value. In such an example, the hashes of large pieces of data may be embedded within transactions, instead of the data itself. Readers will appreciate that, in other embodiments, alternatives to blockchains may be used to facilitate the decentralized storage of information. For example, one alternative to a blockchain that may be used is a blockweave. While conventional blockchains store every transaction to achieve validation, a blockweave permits secure decentralization without the usage of the entire chain, thereby enabling low cost on-chain storage of data. Such blockweaves may utilize a consensus mechanism that is based on proof of access (PoA) and proof of work (PoW).

[0175] The storage systems described above may, either alone or in combination with other computing devices, be used to support in-memory computing applications. In-memory computing involves the storage of information in RAM that is distributed across a cluster of computers. Readers will appreciate that the storage systems described above, especially those that are configurable with customizable amounts of processing resources, storage resources, and memory resources (e.g., those systems in which blades that contain configurable amounts of each type of resource), may be configured in a way so as to provide an infrastructure that can support in-memory computing. Likewise, the storage systems described above may include component parts (e.g., NVDIMMs, 3D crosspoint storage that provide fast random access memory that is persistent) that can actually provide for an improved in-memory computing environment as compared to in-memory computing environments that rely on RAM distributed across dedicated servers.

[0176] In some embodiments, the storage systems described above may be configured to operate as a hybrid in-memory computing environment that includes a universal interface to all storage media (e.g., RAM, flash storage, 3D crosspoint storage). In such embodiments, users may have no knowledge regarding the details of where their data is stored but they can still use the same full, unified API to address data. In such embodiments, the storage system may (in the background) move data to the fastest layer available—including intelligently placing the data in dependence upon various characteristics of the data or in dependence upon some other heuristic. In such an example, the storage systems may even make use of existing products such as Apache Ignite and GridGain to move data between the various storage layers, or the storage systems may make use of custom software to move data between the various storage layers. The storage systems described herein may implement various optimizations to improve the performance of in-memory computing such as, for example, having computations occur as close to the data as possible.

[0177] Readers will further appreciate that in some embodiments, the storage systems described above may be paired with other resources to support the applications described above. For example, one infrastructure could include primary compute in the form of servers and workstations which specialize in using General-purpose computing on graphics processing units (‘GPGPU’) to accelerate deep learning applications that are interconnected into a computation engine to train parameters for deep neural

networks. Each system may have Ethernet external connectivity, InfiniBand external connectivity, some other form of external connectivity, or some combination thereof. In such an example, the GPUs can be grouped for a single large training or used independently to train multiple models. The infrastructure could also include a storage system such as those described above to provide, for example, a scale-out all-flash file or object store through which data can be accessed via high-performance protocols such as NFS, S3, and so on. The infrastructure can also include, for example, redundant top-of-rack Ethernet switches connected to storage and compute via ports in MLAG port channels for redundancy. The infrastructure could also include additional compute in the form of whitebox servers, optionally with GPUs, for data ingestion, pre-processing, and model debugging. Readers will appreciate that additional infrastructures are also possible.

[0178] Readers will appreciate that the storage systems described above, either alone or in coordination with other computing machinery may be configured to support other AI related tools. For example, the storage systems may make use of tools like ONNX or other open neural network exchange formats that make it easier to transfer models written in different AI frameworks. Likewise, the storage systems may be configured to support tools like Amazon's Glue that allow developers to prototype, build, and train deep learning models. In fact, the storage systems described above may be part of a larger platform, such as IBM™ Cloud Private for Data, that includes integrated data science, data engineering and application building services.

[0179] Readers will further appreciate that the storage systems described above may also be deployed as an edge solution. Such an edge solution may be in place to optimize cloud computing systems by performing data processing at the edge of the network, near the source of the data. Edge computing can push applications, data and computing power (i.e., services) away from centralized points to the logical extremes of a network. Through the use of edge solutions such as the storage systems described above, computational tasks may be performed using the compute resources provided by such storage systems, data may be stored using the storage resources of the storage system, and cloud-based services may be accessed through the use of various resources of the storage system (including networking resources). By performing computational tasks on the edge solution, storing data on the edge solution, and generally making use of the edge solution, the consumption of expensive cloud-based resources may be avoided and, in fact, performance improvements may be experienced relative to a heavier reliance on cloud-based resources.

[0180] While many tasks may benefit from the utilization of an edge solution, some particular uses may be especially suited for deployment in such an environment. For example, devices like drones, autonomous cars, robots, and others may require extremely rapid processing—so fast, in fact, that sending data up to a cloud environment and back to receive data processing support may simply be too slow. As an additional example, some IoT devices such as connected video cameras may not be well-suited for the utilization of cloud-based resources as it may be impractical (not only from a privacy perspective, security perspective, or a financial perspective) to send the data to the cloud simply because of the pure volume of data that is involved. As such, many tasks that really rely on data processing, storage, or communi-

cations may be better suited by platforms that include edge solutions such as the storage systems described above.

[0181] The storage systems described above may alone, or in combination with other computing resources, serve as a network edge platform that combines compute resources, storage resources, networking resources, cloud technologies and network virtualization technologies, and so on. As part of the network, the edge may take on characteristics similar to other network facilities, from the customer premise and backhaul aggregation facilities to Points of Presence (PoPs) and regional data centers. Readers will appreciate that network workloads, such as Virtual Network Functions (VNFs) and others, will reside on the network edge platform. Enabled by a combination of containers and virtual machines, the network edge platform may rely on controllers and schedulers that are no longer geographically co-located with the data processing resources. The functions, as microservices, may split into control planes, user and data planes, or even state machines, allowing for independent optimization and scaling techniques to be applied. Such user and data planes may be enabled through increased accelerators, both those residing in server platforms, such as FPGAs and Smart NICs, and through SDN-enabled merchant silicon and programmable ASICs.

[0182] The storage systems described above may also be optimized for use in big data analytics, including being leveraged as part of a composable data analytics pipeline where containerized analytics architectures, for example, make analytics capabilities more composable. Big data analytics may be generally described as the process of examining large and varied data sets to uncover hidden patterns, unknown correlations, market trends, customer preferences and other useful information that can help organizations make more-informed business decisions. As part of that process, semi-structured and unstructured data such as, for example, internet clickstream data, web server logs, social media content, text from customer emails and survey responses, mobile-phone call-detail records, IoT sensor data, and other data may be converted to a structured form.

[0183] The storage systems described above may also support (including implementing as a system interface) applications that perform tasks in response to human speech. For example, the storage systems may support the execution of intelligent personal assistant applications such as, for example, Amazon's Alexa™, Apple Siri™, Google Voice™, Samsung Bixby™, Microsoft Cortana™, and others. While the examples described in the previous sentence make use of voice as input, the storage systems described above may also support chatbots, talkbots, chatterbots, or artificial conversational entities or other applications that are configured to conduct a conversation via auditory or textual methods. Likewise, the storage system may actually execute such an application to enable a user such as a system administrator to interact with the storage system via speech. Such applications are generally capable of voice interaction, music playback, making to-do lists, setting alarms, streaming podcasts, playing audiobooks, and providing weather, traffic, and other real time information, such as news, although in embodiments in accordance with the present disclosure, such applications may be utilized as interfaces to various system management operations.

[0184] The storage systems described above may also implement AI platforms for delivering on the vision of

self-driving storage. Such AI platforms may be configured to deliver global predictive intelligence by collecting and analyzing large amounts of storage system telemetry data points to enable effortless management, analytics and support. In fact, such storage systems may be capable of predicting both capacity and performance, as well as generating intelligent advice on workload deployment, interaction and optimization. Such AI platforms may be configured to scan all incoming storage system telemetry data against a library of issue fingerprints to predict and resolve incidents in real-time, before they impact customer environments, and captures hundreds of variables related to performance that are used to forecast performance load.

[0185] The storage systems described above may support the serialized or simultaneous execution of artificial intelligence applications, machine learning applications, data analytics applications, data transformations, and other tasks that collectively may form an AI ladder. Such an AI ladder may effectively be formed by combining such elements to form a complete data science pipeline, where exist dependencies between elements of the AI ladder. For example, AI may require that some form of machine learning has taken place, machine learning may require that some form of analytics has taken place, analytics may require that some form of data and information architecting has taken place, and so on. As such, each element may be viewed as a rung in an AI ladder that collectively can form a complete and sophisticated AI solution.

[0186] The storage systems described above may also, either alone or in combination with other computing environments, be used to deliver an AI everywhere experience where AI permeates wide and expansive aspects of business and life. For example, AI may play an important role in the delivery of deep learning solutions, deep reinforcement learning solutions, artificial general intelligence solutions, autonomous vehicles, cognitive computing solutions, commercial UAVs or drones, conversational user interfaces, enterprise taxonomies, ontology management solutions, machine learning solutions, smart dust, smart robots, smart workplaces, and many others.

[0187] The storage systems described above may also, either alone or in combination with other computing environments, be used to deliver a wide range of transparently immersive experiences (including those that use digital twins of various “things” such as people, places, processes, systems, and so on) where technology can introduce transparency between people, businesses, and things. Such transparently immersive experiences may be delivered as augmented reality technologies, connected homes, virtual reality technologies, brain-computer interfaces, human augmentation technologies, nanotube electronics, volumetric displays, 4D printing technologies, or others.

[0188] The storage systems described above may also, either alone or in combination with other computing environments, be used to support a wide variety of digital platforms. Such digital platforms can include, for example, 5G wireless systems and platforms, digital twin platforms, edge computing platforms, IoT platforms, quantum computing platforms, serverless PaaS, software-defined security, neuromorphic computing platforms, and so on.

[0189] The storage systems described above may also be part of a multi-cloud environment in which multiple cloud computing and storage services are deployed in a single heterogeneous architecture. In order to facilitate the opera-

tion of such a multi-cloud environment, DevOps tools may be deployed to enable orchestration across clouds. Likewise, continuous development and continuous integration tools may be deployed to standardize processes around continuous integration and delivery, new feature rollout and provisioning cloud workloads. By standardizing these processes, a multi-cloud strategy may be implemented that enables the utilization of the best provider for each workload.

[0190] The storage systems described above may be used as a part of a platform to enable the use of crypto-anchors that may be used to authenticate a product’s origins and contents to ensure that it matches a blockchain record associated with the product. Similarly, as part of a suite of tools to secure data stored on the storage system, the storage systems described above may implement various encryption technologies and schemes, including lattice cryptography. Lattice cryptography can involve constructions of cryptographic primitives that involve lattices, either in the construction itself or in the security proof. Unlike public-key schemes such as the RSA, Diffie-Hellman or Elliptic-Curve cryptosystems, which are easily attacked by a quantum computer, some lattice-based constructions appear to be resistant to attack by both classical and quantum computers.

[0191] A quantum computer is a device that performs quantum computing. Quantum computing is computing using quantum-mechanical phenomena, such as superposition and entanglement. Quantum computers differ from traditional computers that are based on transistors, as such traditional computers require that data be encoded into binary digits (bits), each of which is always in one of two definite states (0 or 1). In contrast to traditional computers, quantum computers use quantum bits, which can be in superpositions of states. A quantum computer maintains a sequence of qubits, where a single qubit can represent a one, a zero, or any quantum superposition of those two qubit states. A pair of qubits can be in any quantum superposition of 4 states, and three qubits in any superposition of 8 states. A quantum computer with n qubits can generally be in an arbitrary superposition of up to 2^n different states simultaneously, whereas a traditional computer can only be in one of these states at any one time. A quantum Turing machine is a theoretical model of such a computer.

[0192] The storage systems described above may also be paired with FPGA-accelerated servers as part of a larger AI or ML infrastructure. Such FPGA-accelerated servers may reside near (e.g., in the same data center) the storage systems described above or even incorporated into an appliance that includes one or more storage systems, one or more FPGA-accelerated servers, networking infrastructure that supports communications between the one or more storage systems and the one or more FPGA-accelerated servers, as well as other hardware and software components. Alternatively, FPGA-accelerated servers may reside within a cloud computing environment that may be used to perform compute-related tasks for AI and ML jobs. Any of the embodiments described above may be used to collectively serve as a FPGA-based AI or ML platform. Readers will appreciate that, in some embodiments of the FPGA-based AI or ML platform, the FPGAs that are contained within the FPGA-accelerated servers may be reconfigured for different types of ML models (e.g., LSTMs, CNNs, GRUs). The ability to reconfigure the FPGAs that are contained within the FPGA-accelerated servers may enable the acceleration of a ML or AI application based on the most optimal numerical preci-

sion and memory model being used. Readers will appreciate that by treating the collection of FPGA-accelerated servers as a pool of FPGAs, any CPU in the data center may utilize the pool of FPGAs as a shared hardware microservice, rather than limiting a server to dedicated accelerators plugged into it.

[0193] The FPGA-accelerated servers and the GPU-accelerated servers described above may implement a model of computing where, rather than keeping a small amount of data in a CPU and running a long stream of instructions over it as occurred in more traditional computing models, the machine learning model and parameters are pinned into the high-bandwidth on-chip memory with lots of data streaming through the high-bandwidth on-chip memory. FPGAs may even be more efficient than GPUs for this computing model, as the FPGAs can be programmed with only the instructions needed to run this kind of computing model.

[0194] The storage systems described above may be configured to provide parallel storage, for example, through the use of a parallel file system such as BeeGFS. Such parallel file systems may include a distributed metadata architecture. For example, the parallel file system may include a plurality of metadata servers across which metadata is distributed, as well as components that include services for clients and storage servers.

[0195] The systems described above can support the execution of a wide array of software applications. Such software applications can be deployed in a variety of ways, including container-based deployment models. Containerized applications may be managed using a variety of tools. For example, containerized applications may be managed using Docker Swarm, Kubernetes, and others. Containerized applications may be used to facilitate a serverless, cloud native computing deployment and management model for software applications. In support of a serverless, cloud native computing deployment and management model for software applications, containers may be used as part of an event handling mechanisms (e.g., AWS Lambdas) such that various events cause a containerized application to be spun up to operate as an event handler.

[0196] The systems described above may be deployed in a variety of ways, including being deployed in ways that support fifth generation ('5G') networks. 5G networks may support substantially faster data communications than previous generations of mobile communications networks and, as a consequence may lead to the disaggregation of data and computing resources as modern massive data centers may become less prominent and may be replaced, for example, by more-local, micro data centers that are close to the mobile-network towers. The systems described above may be included in such local, micro data centers and may be part of or paired to multi-access edge computing ('MEC') systems. Such MEC systems may enable cloud computing capabilities and an IT service environment at the edge of the cellular network. By running applications and performing related processing tasks closer to the cellular customer, network congestion may be reduced and applications may perform better.

[0197] The storage systems described above may also be configured to implement NVMe Zoned Namespaces. Through the use of NVMe Zoned Namespaces, the logical address space of a namespace is divided into zones. Each zone provides a logical block address range that must be written sequentially and explicitly reset before rewriting,

thereby enabling the creation of namespaces that expose the natural boundaries of the device and offload management of internal mapping tables to the host. In order to implement NVMe Zoned Name Spaces ('ZNS'), ZNS SSDs or some other form of zoned block devices may be utilized that expose a namespace logical address space using zones. With the zones aligned to the internal physical properties of the device, several inefficiencies in the placement of data can be eliminated. In such embodiments, each zone may be mapped, for example, to a separate application such that functions like wear levelling and garbage collection could be performed on a per-zone or per-application basis rather than across the entire device. In order to support ZNS, the storage controllers described herein may be configured with to interact with zoned block devices through the usage of, for example, the Linux™ kernel zoned block device interface or other tools.

[0198] The storage systems described above may also be configured to implement zoned storage in other ways such as, for example, through the usage of shingled magnetic recording (SMR) storage devices. In examples where zoned storage is used, device-managed embodiments may be deployed where the storage devices hide this complexity by managing it in the firmware, presenting an interface like any other storage device. Alternatively, zoned storage may be implemented via a host-managed embodiment that depends on the operating system to know how to handle the drive, and only write sequentially to certain regions of the drive. Zoned storage may similarly be implemented using a host-aware embodiment in which a combination of a drive managed and host managed implementation is deployed.

[0199] The storage systems described herein may be used to form a data lake. A data lake may operate as the first place that an organization's data flows to, where such data may be in a raw format. Metadata tagging may be implemented to facilitate searches of data elements in the data lake, especially in embodiments where the data lake contains multiple stores of data, in formats not easily accessible or readable (e.g., unstructured data, semi-structured data, structured data). From the data lake, data may go downstream to a data warehouse where data may be stored in a more processed, packaged, and consumable format. The storage systems described above may also be used to implement such a data warehouse. In addition, a data mart or data hub may allow for data that is even more easily consumed, where the storage systems described above may also be used to provide the underlying storage resources necessary for a data mart or data hub. In embodiments, queries the data lake may require a schema-on-read approach, where data is applied to a plan or schema as it is pulled out of a stored location, rather than as it goes into the stored location.

[0200] The storage systems described herein may also be configured implement a recovery point objective ('RPO'), which may be established by a user, established by an administrator, established as a system default, established as part of a storage class or service that the storage system is participating in the delivery of, or in some other way. A "recovery point objective" is a goal for the maximum time difference between the last update to a source dataset and the last recoverable replicated dataset update that would be correctly recoverable, given a reason to do so, from a continuously or frequently updated copy of the source dataset. An update is correctly recoverable if it properly

takes into account all updates that were processed on the source dataset prior to the last recoverable replicated dataset update.

[0201] In synchronous replication, the RPO would be zero, meaning that under normal operation, all completed updates on the source dataset should be present and correctly recoverable on the copy dataset. In best effort nearly synchronous replication, the RPO can be as low as a few seconds. In snapshot-based replication, the RPO can be roughly calculated as the interval between snapshots plus the time to transfer the modifications between a previous already transferred snapshot and the most recent to-be-replicated snapshot.

[0202] If updates accumulate faster than they are replicated, then an RPO can be missed. If more data to be replicated accumulates between two snapshots, for snapshot-based replication, than can be replicated between taking the snapshot and replicating that snapshot's cumulative updates to the copy, then the RPO can be missed. If, again in snapshot-based replication, data to be replicated accumulates at a faster rate than could be transferred in the time between subsequent snapshots, then replication can start to fall further behind which can extend the miss between the expected recovery point objective and the actual recovery point that is represented by the last correctly replicated update.

[0203] The storage systems described above may also be part of a shared nothing storage cluster. In a shared nothing storage cluster, each node of the cluster has local storage and communicates with other nodes in the cluster through networks, where the storage used by the cluster is (in general) provided only by the storage connected to each individual node. A collection of nodes that are synchronously replicating a dataset may be one example of a shared nothing storage cluster, as each storage system has local storage and communicates to other storage systems through a network, where those storage systems do not (in general) use storage from somewhere else that they share access to through some kind of interconnect. In contrast, some of the storage systems described above are themselves built as a shared-storage cluster, since there are drive shelves that are shared by the paired controllers. Other storage systems described above, however, are built as a shared nothing storage cluster, as all storage is local to a particular node (e.g., a blade) and all communication is through networks that link the compute nodes together.

[0204] In other embodiments, other forms of a shared nothing storage cluster can include embodiments where any node in the cluster has a local copy of all storage they need, and where data is mirrored through a synchronous style of replication to other nodes in the cluster either to ensure that the data isn't lost or because other nodes are also using that storage. In such an embodiment, if a new cluster node needs some data, that data can be copied to the new node from other nodes that have copies of the data.

[0205] In some embodiments, mirror-copy-based shared storage clusters may store multiple copies of all the cluster's stored data, with each subset of data replicated to a particular set of nodes, and different subsets of data replicated to different sets of nodes. In some variations, embodiments may store all of the cluster's stored data in all nodes, whereas in other variations nodes may be divided up such

that a first set of nodes will all store the same set of data and a second, different set of nodes will all store a different set of data.

[0206] Readers will appreciate that RAFT-based databases (e.g., etcd) may operate like shared-nothing storage clusters where all RAFT nodes store all data. The amount of data stored in a RAFT cluster, however, may be limited so that extra copies don't consume too much storage. A container server cluster might also be able to replicate all data to all cluster nodes, presuming the containers don't tend to be too large and their bulk data (the data manipulated by the applications that run in the containers) is stored elsewhere such as in an S3 cluster or an external file server. In such an example, the container storage may be provided by the cluster directly through its shared-nothing storage model, with those containers providing the images that form the execution environment for parts of an application or service.

[0207] For further explanation, FIG. 3D illustrates an exemplary computing device 350 that may be specifically configured to perform one or more of the processes described herein. As shown in FIG. 3D, computing device 350 may include a communication interface 352, a processor 354, a storage device 356, and an input/output ("I/O") module 358 communicatively connected one to another via a communication infrastructure 360. While an exemplary computing device 350 is shown in FIG. 3D, the components illustrated in FIG. 3D are not intended to be limiting. Additional or alternative components may be used in other embodiments. Components of computing device 350 shown in FIG. 3D will now be described in additional detail.

[0208] Communication interface 352 may be configured to communicate with one or more computing devices. Examples of communication interface 352 include, without limitation, a wired network interface (such as a network interface card), a wireless network interface (such as a wireless network interface card), a modem, an audio/video connection, and any other suitable interface.

[0209] Processor 354 generally represents any type or form of processing unit capable of processing data and/or interpreting, executing, and/or directing execution of one or more of the instructions, processes, and/or operations described herein. Processor 354 may perform operations by executing computer-executable instructions 362 (e.g., an application, software, code, and/or other executable data instance) stored in storage device 356.

[0210] Storage device 356 may include one or more data storage media, devices, or configurations and may employ any type, form, and combination of data storage media and/or device. For example, storage device 356 may include, but is not limited to, any combination of the non-volatile media and/or volatile media described herein. Electronic data, including data described herein, may be temporarily and/or permanently stored in storage device 356. For example, data representative of computer-executable instructions 362 configured to direct processor 354 to perform any of the operations described herein may be stored within storage device 356. In some examples, data may be arranged in one or more databases residing within storage device 356.

[0211] I/O module 358 may include one or more I/O modules configured to receive user input and provide user output. I/O module 358 may include any hardware, firmware, software, or combination thereof supportive of input and output capabilities. For example, I/O module 358 may

include hardware and/or software for capturing user input, including, but not limited to, a keyboard or keypad, a touchscreen component (e.g., touchscreen display), a receiver (e.g., an RF or infrared receiver), motion sensors, and/or one or more input buttons.

[0212] I/O module 358 may include one or more devices for presenting output to a user, including, but not limited to, a graphics engine, a display (e.g., a display screen), one or more output drivers (e.g., display drivers), one or more audio speakers, and one or more audio drivers. In certain embodiments, I/O module 358 is configured to provide graphical data to a display for presentation to a user. The graphical data may be representative of one or more graphical user interfaces and/or any other graphical content as may serve a particular implementation. In some examples, any of the systems, computing devices, and/or other components described herein may be implemented by computing device 350.

[0213] For further explanation, FIG. 3E illustrates an example of a fleet of storage systems 376 for providing storage services (also referred to herein as ‘data services’). The fleet of storage systems 376 depicted in FIG. 3 includes a plurality of storage systems 374a, 374b, 374c, 374d, 374n that may each be similar to the storage systems described herein. The storage systems 374a, 374b, 374c, 374d, 374n in the fleet of storage systems 376 may be embodied as identical storage systems or as different types of storage systems. For example, two of the storage systems 374a, 374n depicted in FIG. 3E are depicted as being cloud-based storage systems, as the resources that collectively form each of the storage systems 374a, 374n are provided by distinct cloud services providers 370, 372. For example, the first cloud services provider 370 may be Amazon AWS™ whereas the second cloud services provider 372 is Microsoft Azure™, although in other embodiments one or more public clouds, private clouds, or combinations thereof may be used to provide the underlying resources that are used to form a particular storage system in the fleet of storage systems 376.

[0214] The example depicted in FIG. 3E includes an edge management service 382 for delivering storage services in accordance with some embodiments of the present disclosure. The storage services (also referred to herein as ‘data services’) that are delivered may include, for example, services to provide a certain amount of storage to a consumer, services to provide storage to a consumer in accordance with a predetermined service level agreement, services to provide storage to a consumer in accordance with predetermined regulatory requirements, and many others.

[0215] The edge management service 382 depicted in FIG. 3E may be embodied, for example, as one or more modules of computer program instructions executing on computer hardware such as one or more computer processors. Alternatively, the edge management service 382 may be embodied as one or more modules of computer program instructions executing on a virtualized execution environment such as one or more virtual machines, in one or more containers, or in some other way. In other embodiments, the edge management service 382 may be embodied as a combination of the embodiments described above, including embodiments where the one or more modules of computer program instructions that are included in the edge management service 382 are distributed across multiple physical or virtual execution environments.

[0216] The edge management service 382 may operate as a gateway for providing storage services to storage consumers, where the storage services leverage storage offered by one or more storage systems 374a, 374b, 374c, 374d, 374n. For example, the edge management service 382 may be configured to provide storage services to host devices 378a, 378b, 378c, 378d, 378n that are executing one or more applications that consume the storage services. In such an example, the edge management service 382 may operate as a gateway between the host devices 378a, 378b, 378c, 378d, 378n and the storage systems 374a, 374b, 374c, 374d, 374n, rather than requiring that the host devices 378a, 378b, 378c, 378d, 378n directly access the storage systems 374a, 374b, 374c, 374d, 374n.

[0217] The edge management service 382 of FIG. 3E exposes a storage services module 380 to the host devices 378a, 378b, 378c, 378d, 378n of FIG. 3E, although in other embodiments the edge management service 382 may expose the storage services module 380 to other consumers of the various storage services. The various storage services may be presented to consumers via one or more user interfaces, via one or more APIs, or through some other mechanism provided by the storage services module 380. As such, the storage services module 380 depicted in FIG. 3E may be embodied as one or more modules of computer program instructions executing on physical hardware, on a virtualized execution environment, or combinations thereof, where executing such modules causes enables a consumer of storage services to be offered, select, and access the various storage services.

[0218] The edge management service 382 of FIG. 3E also includes a system management services module 384. The system management services module 384 of FIG. 3E includes one or more modules of computer program instructions that, when executed, perform various operations in coordination with the storage systems 374a, 374b, 374c, 374d, 374n to provide storage services to the host devices 378a, 378b, 378c, 378d, 378n. The system management services module 384 may be configured, for example, to perform tasks such as provisioning storage resources from the storage systems 374a, 374b, 374c, 374d, 374n via one or more APIs exposed by the storage systems 374a, 374b, 374c, 374d, 374n, migrating datasets or workloads amongst the storage systems 374a, 374b, 374c, 374d, 374n via one or more APIs exposed by the storage systems 374a, 374b, 374c, 374d, 374n, setting one or more tunable parameters (i.e., one or more configurable settings) on the storage systems 374a, 374b, 374c, 374d, 374n via one or more APIs exposed by the storage systems 374a, 374b, 374c, 374d, 374n, and so on. For example, many of the services described below relate to embodiments where the storage systems 374a, 374b, 374c, 374d, 374n are configured to operate in some way. In such examples, the system management services module 384 may be responsible for using APIs (or some other mechanism) provided by the storage systems 374a, 374b, 374c, 374d, 374n to configure the storage systems 374a, 374b, 374c, 374d, 374n to operate in the ways described below.

[0219] In addition to configuring the storage systems 374a, 374b, 374c, 374d, 374n, the edge management service 382 itself may be configured to perform various tasks required to provide the various storage services. Consider an example in which the storage service includes a service that, when selected and applied, causes personally identifiable

information ('PII') contained in a dataset to be obfuscated when the dataset is accessed. In such an example, the storage systems 374a, 374b, 374c, 374d, 374n may be configured to obfuscate PII when servicing read requests directed to the dataset. Alternatively, the storage systems 374a, 374b, 374c, 374d, 374n may service reads by returning data that includes the PII, but the edge management service 382 itself may obfuscate the PII as the data is passed through the edge management service 382 on its way from the storage systems 374a, 374b, 374c, 374d, 374n to the host devices 378a, 378b, 378c, 378d, 378n.

[0220] The storage systems 374a, 374b, 374c, 374d, 374n depicted in FIG. 3E may be embodied as one or more of the storage systems described above with reference to FIGS. 1A-3D, including variations thereof. In fact, the storage systems 374a, 374b, 374c, 374d, 374n may serve as a pool of storage resources where the individual components in that pool have different performance characteristics, different storage characteristics, and so on. For example, one of the storage systems 374a may be a cloud-based storage system, another storage system 374b may be a storage system that provides block storage, another storage system 374c may be a storage system that provides file storage, another storage system 374d may be a relatively high-performance storage system while another storage system 374n may be a relatively low-performance storage system, and so on. In alternative embodiments, only a single storage system may be present.

[0221] The storage systems 374a, 374b, 374c, 374d, 374n depicted in FIG. 3E may also be organized into different failure domains so that the failure of one storage system 374a should be totally unrelated to the failure of another storage system 374b. For example, each of the storage systems may receive power from independent power systems, each of the storage systems may be coupled for data communications over independent data communications networks, and so on. Furthermore, the storage systems in a first failure domain may be accessed via a first gateway whereas storage systems in a second failure domain may be accessed via a second gateway. For example, the first gateway may be a first instance of the edge management service 382 and the second gateway may be a second instance of the edge management service 382, including embodiments where each instance is distinct, or each instance is part of a distributed edge management service 382.

[0222] As an illustrative example of available storage services, storage services may be presented to a user that are associated with different levels of data protection. For example, storage services may be presented to the user that, when selected and enforced, guarantee the user that data associated with that user will be protected such that various recovery point objectives ('RPO') can be guaranteed. A first available storage service may ensure, for example, that some dataset associated with the user will be protected such that any data that is more than 5 seconds old can be recovered in the event of a failure of the primary data store whereas a second available storage service may ensure that the dataset that is associated with the user will be protected such that any data that is more than 5 minutes old can be recovered in the event of a failure of the primary data store.

[0223] An additional example of storage services that may be presented to a user, selected by a user, and ultimately applied to a dataset associated with the user can include one

or more data compliance services. Such data compliance services may be embodied, for example, as services that may be provided to consumers (i.e., a user) the data compliance services to ensure that the user's datasets are managed in a way to adhere to various regulatory requirements. For example, one or more data compliance services may be offered to a user to ensure that the user's datasets are managed in a way so as to adhere to the General Data Protection Regulation ('GDPR'), one or data compliance services may be offered to a user to ensure that the user's datasets are managed in a way so as to adhere to the Sarbanes—Oxley Act of 2002 ('SOX'), or one or more data compliance services may be offered to a user to ensure that the user's datasets are managed in a way so as to adhere to some other regulatory act. In addition, the one or more data compliance services may be offered to a user to ensure that the user's datasets are managed in a way so as to adhere to some non-governmental guidance (e.g., to adhere to best practices for auditing purposes), the one or more data compliance services may be offered to a user to ensure that the user's datasets are managed in a way so as to adhere to a particular clients or organizations requirements, and so on.

[0224] Consider an example in which a particular data compliance service is designed to ensure that a user's datasets are managed in a way so as to adhere to the requirements set forth in the GDPR. While a listing of all requirements of the GDPR can be found in the regulation itself, for the purposes of illustration, an example requirement set forth in the GDPR requires that pseudonymization processes must be applied to stored data in order to transform personal data in such a way that the resulting data cannot be attributed to a specific data subject without the use of additional information. For example, data encryption techniques can be applied to render the original data unintelligible, and such data encryption techniques cannot be reversed without access to the correct decryption key. As such, the GDPR may require that the decryption key be kept separately from the pseudonymised data. One particular data compliance service may be offered to ensure adherence to the requirements set forth in this paragraph.

[0225] In order to provide this particular data compliance service, the data compliance service may be presented to a user (e.g., via a GUI) and selected by the user. In response to receiving the selection of the particular data compliance service, one or more storage services policies may be applied to a dataset associated with the user to carry out the particular data compliance service. For example, a storage services policy may be applied requiring that the dataset be encrypted prior to be stored in a storage system, prior to being stored in a cloud environment, or prior to being stored elsewhere. In order to enforce this policy, a requirement may be enforced not only requiring that the dataset be encrypted when stored, but a requirement may be put in place requiring that the dataset be encrypted prior to transmitting the dataset (e.g., sending the dataset to another party). In such an example, a storage services policy may also be put in place requiring that any encryption keys used to encrypt the dataset are not stored on the same system that stores the dataset itself. Readers will appreciate that many other forms of data compliance services may be offered and implemented in accordance with embodiments of the present disclosure.

[0226] The storage systems 374a, 374b, 374c, 374d, 374n in the fleet of storage systems 376 may be managed collec-

tively, for example, by one or more fleet management modules. The fleet management modules may be part of or separate from the system management services module 384 depicted in FIG. 3E. The fleet management modules may perform tasks such as monitoring the health of each storage system in the fleet, initiating updates or upgrades on one or more storage systems in the fleet, migrating workloads for loading balancing or other performance purposes, and many other tasks. As such, and for many other reasons, the storage systems 374a, 374b, 374c, 374d, 374n may be coupled to each other via one or more data communications links in order to exchange data between the storage systems 374a, 374b, 374c, 374d, 374n.

[0227] The storage systems described herein may support various forms of data replication. For example, two or more of the storage systems may synchronously replicate a dataset between each other. In synchronous replication, distinct copies of a particular dataset may be maintained by multiple storage systems, but all accesses (e.g., a read) of the dataset should yield consistent results regardless of which storage system the access was directed to. For example, a read directed to any of the storage systems that are synchronously replicating the dataset should return identical results. As such, while updates to the version of the dataset need not occur at exactly the same time, precautions must be taken to ensure consistent accesses to the dataset. For example, if an update (e.g., a write) that is directed to the dataset is received by a first storage system, the update may only be acknowledged as being completed if all storage systems that are synchronously replicating the dataset have applied the update to their copies of the dataset. In such an example, synchronous replication may be carried out through the use of I/O forwarding (e.g., a write received at a first storage system is forwarded to a second storage system), communications between the storage systems (e.g., each storage system indicating that it has completed the update), or in other ways.

[0228] In other embodiments, a dataset may be replicated through the use of checkpoints. In checkpoint-based replication (also referred to as ‘nearly synchronous replication’), a set of updates to a dataset (e.g., one or more write operations directed to the dataset) may occur between different checkpoints, such that a dataset has been updated to a specific checkpoint only if all updates to the dataset prior to the specific checkpoint have been completed. Consider an example in which a first storage system stores a live copy of a dataset that is being accessed by users of the dataset. In this example, assume that the dataset is being replicated from the first storage system to a second storage system using checkpoint-based replication. For example, the first storage system may send a first checkpoint (at time $t=0$) to the second storage system, followed by a first set of updates to the dataset, followed by a second checkpoint (at time $t=1$), followed by a second set of updates to the dataset, followed by a third checkpoint (at time $t=2$). In such an example, if the second storage system has performed all updates in the first set of updates but has not yet performed all updates in the second set of updates, the copy of the dataset that is stored on the second storage system may be up-to-date until the second checkpoint. Alternatively, if the second storage system has performed all updates in both the first set of updates and the second set of updates, the copy of the dataset that is stored on the second storage system may be up-to-date until the third checkpoint. Readers will appreciate that various

types of checkpoints may be used (e.g., metadata only checkpoints), checkpoints may be spread out based on a variety of factors (e.g., time, number of operations, an RPO setting), and so on.

[0229] In other embodiments, a dataset may be replicated through snapshot-based replication (also referred to as ‘asynchronous replication’). In snapshot-based replication, snapshots of a dataset may be sent from a replication source such as a first storage system to a replication target such as a second storage system. In such an embodiment, each snapshot may include the entire dataset or a subset of the dataset such as, for example, only the portions of the dataset that have changed since the last snapshot was sent from the replication source to the replication target. Readers will appreciate that snapshots may be sent on-demand, based on a policy that takes a variety of factors into consideration (e.g., time, number of operations, an RPO setting), or in some other way.

[0230] The storage systems described above may, either alone or in combination, be configured to serve as a continuous data protection store. A continuous data protection store is a feature of a storage system that records updates to a dataset in such a way that consistent images of prior contents of the dataset can be accessed with a low time granularity (often on the order of seconds, or even less), and stretching back for a reasonable period of time (often hours or days). These allow access to very recent consistent points in time for the dataset, and also allow access to access to points in time for a dataset that might have just preceded some event that, for example, caused parts of the dataset to be corrupted or otherwise lost, while retaining close to the maximum number of updates that preceded that event. Conceptually, they are like a sequence of snapshots of a dataset taken very frequently and kept for a long period of time, though continuous data protection stores are often implemented quite differently from snapshots. A storage system implementing a data continuous data protection store may further provide a means of accessing these points in time, accessing one or more of these points in time as snapshots or as cloned copies, or reverting the dataset back to one of those recorded points in time.

[0231] Over time, to reduce overhead, some points in the time held in a continuous data protection store can be merged with other nearby points in time, essentially deleting some of these points in time from the store. This can reduce the capacity needed to store updates. It may also be possible to convert a limited number of these points in time into longer duration snapshots. For example, such a store might keep a low granularity sequence of points in time stretching back a few hours from the present, with some points in time merged or deleted to reduce overhead for up to an additional day. Stretching back in the past further than that, some of these points in time could be converted to snapshots representing consistent point-in-time images from only every few hours.

[0232] Although some embodiments are described largely in the context of a storage system, readers of skill in the art will recognize that embodiments of the present disclosure may also take the form of a computer program product disposed upon computer readable storage media for use with any suitable processing system. Such computer readable storage media may be any storage medium for machine-readable information, including magnetic media, optical media, solid-state media, or other suitable media. Examples

of such media include magnetic disks in hard drives or diskettes, compact disks for optical drives, magnetic tape, and others as will occur to those of skill in the art. Persons skilled in the art will immediately recognize that any computer system having suitable programming means will be capable of executing the steps described herein as embodied in a computer program product. Persons skilled in the art will recognize also that, although some of the embodiments described in this specification are oriented to software installed and executing on computer hardware, nevertheless, alternative embodiments implemented as firmware or as hardware are well within the scope of the present disclosure.

[0233] In some examples, a non-transitory computer-readable medium storing computer-readable instructions may be provided in accordance with the principles described herein. The instructions, when executed by a processor of a computing device, may direct the processor and/or computing device to perform one or more operations, including one or more of the operations described herein. Such instructions may be stored and/or transmitted using any of a variety of known computer-readable media.

[0234] A non-transitory computer-readable medium as referred to herein may include any non-transitory storage medium that participates in providing data (e.g., instructions) that may be read and/or executed by a computing device (e.g., by a processor of a computing device). For example, a non-transitory computer-readable medium may include, but is not limited to, any combination of non-volatile storage media and/or volatile storage media. Exemplary non-volatile storage media include, but are not limited to, read-only memory, flash memory, a solid-state drive, a magnetic storage device (e.g., a hard disk, a floppy disk, magnetic tape, etc.), ferroelectric random-access memory (“RAM”), and an optical disc (e.g., a compact disc, a digital video disc, a Blu-ray disc, etc.). Exemplary volatile storage media include, but are not limited to, RAM (e.g., dynamic RAM).

[0235] While particular combinations of various functions and features of the one or more embodiments are expressly described herein, other combinations of these features and functions are likewise possible. The present disclosure is not limited by the particular examples disclosed herein and expressly incorporates these other combinations.

[0236] In some embodiments, one or more storage systems or one or more elements of storage systems (e.g., features, services, operations, components, etc. of storage systems), such as any of the illustrative storage systems or storage system elements described herein, may be implemented in one or more container systems. A container system may include any system that supports execution of one or more containerized applications. A containerized application may include a software application and an entire runtime environment of the software application bundled into a single image. For example, the containerized application may include source code of the software application and various dependencies, libraries, and/or other components used by the software application to operate. As a result, the containerized application may be abstracted away from a host operating system as a lightweight and portable package, where the containerized application may be uniformly deployed and consistently executed on different computing environments that use different operating systems or different infrastructures. In some embodiments, a containerized application shares a kernel with a host com-

puter system and executes as an isolated user space within an operating system of the host computer system. When executed, a containerized application may provide one or more containerized workloads and/or services.

[0237] The container system may include and/or utilize a cluster of nodes. For example, the container system may be configured to manage deployment and execution of containerized applications on one or more nodes in a cluster. The containerized applications may utilize resources of the nodes, such as processing and/or storage resources provided and/or accessed by the nodes. The storage resources may include any of the illustrative storage resources described herein and may include on-node resources (i.e., storage resources that are locally accessible at the node), off-node resources (i.e., storage resources that are not locally accessible at the node but, for example, may be locally accessible at other nodes), or both on-node and off-node resources.

[0238] In some embodiments, the container system may include a container orchestration system (which may also be referred to as a container orchestrator, a container orchestration platform, etc.) configured to automatically deploy, scale, and manage containerized applications. In some embodiments, the container system may include a storage management system configured to provision and manage storage resources (e.g., virtual volumes) for use by cluster nodes and/or containerized applications.

[0239] In some embodiments, the storage management system may be configured to provide and use different types of network paths to send and receive storage operations and/or data within the container system (e.g., between cluster nodes of the container system). For example, the storage management system may include a container storage system configured to provide a first type of network path through a container system communication network (e.g., through a user space of the container storage system) and a second type of network path through a storage area network. The container storage system may be implemented, for instance, as a containerized application (e.g., a privileged containerized application that is endowed with privileges not afforded to other containerized applications so as to allow the container storage system to carry out storage-related tasks in kernel space of the node), as a non-containerized application executing on the node, or in any other suitable form.

[0240] In certain embodiments, storage operations (e.g., I/O operations) may be dynamically and intelligently routed among a plurality of network paths of different types within a container system using the container storage system. For example, in the container system, a containerized application may access a storage resource over a plurality of network paths of different types. Each network path of the plurality of network paths may be distinct from another network path among the plurality of network paths. The container storage system may intelligently select a network path, from among the plurality of network paths, for routing one or more storage operations (e.g., from the containerized application to the storage resource) based on conditions in the container system, such as network conditions of the plurality of network paths. In addition, the container storage system may intelligently select another network path, from among the plurality of network paths, for routing one or more additional storage operations (e.g., from the containerized application to the storage resource) based on changed

conditions in the container system, such as changed network conditions of the plurality of network paths.

[0241] By providing different types of network paths and intelligently routing storage operations over the different types of network paths in the container system in any of the ways described herein, a first type of network path may be used for certain storage operations when certain conditions exist, and a second type of network path, which is different from the first type of network path, may be used for certain storage operations when other conditions exist. In certain implementations, one of the types of network paths (e.g., a network path through a storage area network) may be relatively faster and/or more efficient than another of the types of network paths (e.g., a network path through a user space of the container storage system) and may be intelligently selected for routing of storage operations when defined conditions exist. Examples of intelligent routing of storage operations across different types of network paths will now be described in reference to the figures.

[0242] FIGS. 4A and 4B illustrate an example container system 400 in which intelligent storage path routing is implemented. In this example, the container system 400 may include a container storage system 402 that dynamically routes storage operations among available network paths based on conditions in the container system 400, such as network conditions of the network paths, as described herein. For example, the container storage system 402 may be configured to route storage operations from containerized applications 404-1 through 404-L to storage resources in one or more storage pools 406 via one or more network paths that are dynamically and intelligently selected from a plurality of available network paths, as described herein.

[0243] The container system 400 may include or be implemented by one or more container orchestration systems, including Kubernetes™, Mesos™, Docker Swarm™, among others. The container orchestration system may manage the container system 400 running on a cluster 408 through services implemented by a control node, depicted as 410.

[0244] A control plane of the container system 400 may implement services that include: deploying applications via a controller 412, monitoring applications via the controller 412, providing an interface via an API server 414, and scheduling deployments via scheduler 416. In some examples, controller 412, scheduler 416, API server 414, and container storage system 402 are implemented by a single node, node 410. In other examples, for resiliency, the control plane may be implemented by multiple, redundant nodes, where if a node that is providing management services for the container system 400 fails, then another, redundant node may provide management services for the cluster 408.

[0245] A data plane of the container system 400 may include a set of nodes that provides container runtimes for executing containerized applications. An individual node within the cluster 408 may execute a container runtime, such as Docker™, and execute a container manager, or node agent, such as a kubelet in Kubernetes (not depicted) that communicates with the control plane via a network proxy, such as an agent 418. The agent 418 may route network traffic to and from containers using, for example, Internet Protocol (IP) port numbers. For example, a containerized application may request a storage class from the control plane, where the request is handled by the container man-

ager, and the container manager communicates the request to the control plane using the agent 418.

[0246] Cluster 408 may include a set of nodes that run containerized applications. A node may be a virtual or physical machine. A virtual machine is an emulation of a computer system, including a kernel, operating system, and namespace, where the emulated computer system may be managed by a virtual machine manager that executes on a host computer system. A virtual machine may be implemented by a compute instance within a cloud-based computing environment or by a physical machine. A process or application executing within a virtual machine does not communicate directly with a host computer system or host kernel, rather a process within a virtual machine communicates with the kernel and operating system of the emulated computer system, and kernel and operating system operations of the emulated computer system are handled by the virtual machine manager, which communicates with a host computer system.

[0247] By contrast, a containerized application shares a kernel with a host computer system and executes as an isolated user space within an operating system of a host system. A host system may be a node. In some embodiments, a containerized application may include a software application and an entire runtime environment of the software application bundled into a single image. For example, as described above, the containerized application may include source code of the software application and various dependencies, libraries, or other components used by the software application to operate. As a result, the containerized application may be abstracted away from a host operating system as a lightweight and portable package, where the containerized application may be uniformly deployed and consistently executed on different computing environments that use different operating systems or different infrastructures.

[0248] The container storage system 402 may orchestrate storage resources to provide storage to the container system 400. For example, the container storage system 402 may provide persistent storage to the containerized applications 404 using the storage pool 406. The container storage system 402 may be deployed as a containerized application by a container orchestration system. For example, the container storage system 402 application may be deployed within cluster 408 and perform management functions for providing storage to the containerized applications 404. Management functions may include determining one or more storage pools from available storage resources, provisioning virtual volumes on one or more nodes, replicating data, and/or handling storage operations. The storage pool 406 may include storage resources from one or more local or remote sources, where the storage resources may be different types of storage.

[0249] The container storage system 402 may also be deployed on a set of nodes for which persistent storage may be provided by the container orchestration system. In some examples, the container storage system 402 may be deployed on all nodes in a cluster 408 using, for example, a DaemonSet. In this example, nodes 420-1 through 420-N provide a container runtime where container storage system 402 executes. In other examples, some, but not all nodes in a cluster may execute the container storage system 402.

[0250] The container storage system 402 may handle storage on a node and communicate with the control plane of container system 400, to provide dynamic volumes,

including persistent volumes. A persistent volume may be mounted on a node as a virtual volume, such as virtual volumes **422-1** and **422-P**. After a virtual volume **422** is mounted, containerized applications may request and use storage provided by the virtual volume **422**. In some implementations, the container storage system **402** may install a driver on a kernel of a node, where the driver handles storage operations directed to the virtual volume. The driver may receive a storage operation directed to a virtual volume, and in response, the driver may perform the storage operation on one or more storage resources within the storage pool **406**.

[0251] The container storage system **402** may, in response to being deployed, determine available storage resources. For example, storage resources **424-1** through **424-M** may include local storage, remote storage, or both local and remote storage. The storage resources **424-1** through **424-M** may include any type(s) and/or configuration(s) of storage resources (e.g., any of the illustrative storage resources described above), and the container storage system **402** may be configured to determine the available storage resources in any suitable way, including based on a configuration file. For example, a configuration file may specify account and authentication information for cloud-based object storage **348** or for a cloud-based storage system **318**. The container storage system **402** may also determine availability of one or more storage devices **356** or one or more storage arrays **102**. An aggregate amount of storage from one or more of storage device(s) **356**, storage array(s) **102**, cloud-based storage system(s) **318**, edge management services **382**, cloud-based object storage **348**, or any other storage resources, or any combination or sub-combination of such storage resources may be used to provide the storage pool **406**. The storage pool **406** is used to provision storage for the one or more virtual volumes **422** mounted on one or more of the nodes **420** within cluster **408**.

[0252] In some implementations, the container storage system **402** may create multiple storage pools. For example, the container storage system **402** may aggregate storage resources of a same type into an individual storage pool. In this example, a storage type may be one of: a storage device **356**, a storage array **102**, a cloud-based storage system **318**, storage via an edge management service **382**, or a cloud-based object storage **348**.

[0253] The container storage system **402** may execute within the cluster **408**, where instances of the container storage system **402** may operate on different nodes within the cluster **408**. In some examples, the container storage system **402** may operate with the container orchestration system of the container system **400** to handle storage operations, mount virtual volumes to provide storage to a node, aggregate available storage into a storage pool **406**, provision storage for a virtual volume from a storage pool **406**, generate backup data, among other storage system operations. In some examples, the container storage system **402** may provide storage services across multiple clusters operating in distinct computing environments. For example, other storage system operations may include storage system operations described above with respect to FIGS. 1-3. Persistent storage provided by the container storage system **402** may be used to implement stateful and/or resilient containerized applications.

[0254] FIG. 4B illustrates an example configuration of the container system **400** in which the container storage system **402** may include modules **426** that provide control plane and

data plane services for storage operations provided by the container storage system **402**. The control plane module **426-1** and the data plane module **426-2** may communicate over the cluster **408** using various communication protocols, such as gRPC, HTTP, among others.

[0255] In some implementations, the control plane module **426-1** may handle one or more of: clustering, provisioning, or scheduling interactions with the container system **400**. For example, various ones of the nodes **420** may be organized into clusters, where an individual cluster may share one or more storage pools. The control plane module **426-1** may store membership information of nodes, and configuration for one or more volumes of a storage pool, where the membership and configuration information may be stored in a database, such as a key-value database.

[0256] The control plane module **426-1** may receive metrics from each node to use to determine scheduling of storage operations. In some implementations, the metrics may include one or more of: node health information, I/O loads, CPU loads, memory loads, among others. The control plane module **426-1** may determine provisioning of storage and data placement within a storage pool based on one or more of: numbers of replicas specified by a configuration setting, storage pool categories or types, specified storage pools, or topology constraints.

[0257] The data plane module **426-2** may handle storage operations and communication with storage resources(s) **424**. For example, the data plane module **426-2** may include an I/O queue, a write-through cache, a timestamp log of storage operations, an I/O dispatcher, device queues, and logic for replicating stored data.

[0258] In the configuration of the container system **400** shown in FIG. 4B, examples of storage resources **424** that may be included in the storage pool **406** are illustrated. For example, storage pool **406** may include one or more storage resources **424** (e.g., storage resources **424-1** through **424-M**) that comprise the storage pool **406**. The storage resources **424** may include one or more of local and/or remote storage devices such as storage device(s) **356**, storage device(s) **256**, storage array(s) **102**, cloud-based object storage **348**, and/or edge management service **382**. Such storage devices may include one or more of: solid-state non-volatile memory, non-mechanical storage device(s) such as solid-state non-volatile memory, mechanical storage devices such as hard-disk drives, and/or any other devices that provide resources for data storage. The one or more storage arrays **102** may be implemented as described with reference to FIG. 1. Cloud-based object storage **348** may be implemented as described with respect to FIGS. 3A-3C. An edge management service **382** may provide storage services as described with respect to FIG. 3E.

[0259] In some embodiments, container storage system **402** may be configured to perform one or more operations to dynamically determine storage path routing within the container system **400**. Examples of such storage path routing will now be described in more detail. It is contemplated that one or more of the illustrative storage path routing operations described herein may be implemented in any suitable way in any suitable container system.

[0260] FIG. 5 illustrates an example configuration for storage path routing within a container system, in accordance with some embodiments of the present disclosure. The container storage system **402** may be deployed within the container system **400**. The container storage system **402** may

include a storage path routing module **500**. In FIG. **5**, the storage path routing module **500** is illustrated as being included in an instance of the container storage system **402** implemented by the control node **410**. This is illustrative only. One or more instances of the storage path routing module **500** may be implemented by one or more of the nodes of the cluster **408**, including the control node **410** and/or nodes **420**, in any suitable configuration.

[0261] The storage path routing module **500** may receive a storage operation directed to a virtual volume **422**, and in response, the storage path routing module **500** may route the storage operation to a host storage resource. The host storage resource may be a storage resource associated with the virtual volume **422**.

[0262] The container storage system **402** (e.g., the storage path routing module **500** of the container storage system **402**) may dynamically route the storage operation based on one or more conditions within the container system **400**. For example, based on network conditions, the container storage system **402** may select a network path, from among a plurality of available network paths, that satisfies one or more metrics, such as network latency, cost, among others. The selected network path may serve as a storage path for routing the storage operation to the host storage resource.

[0263] To illustrate, the cluster **408** of the container system **400** may include a plurality of network paths of different types, including a first network path of a first type and a second network path of a second type. In some implementations, the first network path may be a path through a storage area network **502**, and the second network path may be a path through a local area network **504**. The network paths may be used as storage paths between containerized applications **506**, **508** and one or more storage resources. The storage paths may be used to route storage operations from containerized applications **506**, **508** to the one or more storage resources. A storage resource may be any storage resource included within a storage pool **406** that is provided by the container storage system **402**.

[0264] In some implementations, a storage path through the storage area network **502** may allow storage operations to be routed over the storage path via the kernel space(s) of one or more nodes of the cluster **408**, without utilizing a container system communication network such as a user space of a containerized application. To illustrate, the container storage system **402** may handle storage on the node **420-N** and communicate with the control plane of container system **400** to provide dynamic volumes, including persistent volumes. A persistent volume may be mounted on the node **420-N** as the virtual volume **422-P**. After the virtual volume **422-P** is mounted, the containerized application **508** may request and use storage provided by the virtual volume **422-P**. In some implementations, the container storage system **402** may install a driver on a kernel of the node **420-N**, where the driver handles storage operations directed to the virtual volume **422-P**. The driver may include the storage path routing module **500**.

[0265] In some implementations, network paths over the storage area network **502** may generally be faster than network paths over the local area network **504**. Communication over the storage area network **502** may be performed via kernel space of a node **420** in the container system **400**. The storage area network **502** may be an independent network relative to the local area network **504**. A storage operation may be received by a driver executing within the

kernel of an operating system for a node **420**. As noted above, the driver may include the storage path routing module **500**. In some implementations, the storage area network **502** may generally have higher performance and bandwidth than the local area network **504**, and the local area network **504** may generally be more resilient than a local area network **502**.

[0266] The storage path routing module **500** may determine whether to route the storage operation over the storage area network **502**, over the local area network **504**, or another network among the plurality of networks. If the storage path routing module **500** determines to route the storage operation over the storage area network **502**, then the storage operation is communicated to a storage resource via a network path of the storage area network **502** using a storage area network protocol.

[0267] The storage area network **502** may be implemented with a variety of data communications fabrics, devices, and protocols. For example, fabrics for storage area network **502** may include Fibre Channel, Ethernet, Infiniband, Serial Attached Small Computer System Interface (SAS), or the like. Data communications protocols for use with storage area network **502** may include Advanced Technology Attachment (ATA), Fibre Channel Protocol, Small Computer System Interface (SCSI), Internet Small Computer System Interface (iSCSI), HyperSCSI, Non-Volatile Memory Express (NVMe) over Fabrics, among others.

[0268] In contrast, communication over the local area network **504** may be performed via the user space of the container system **400**. For example, a node **420** may have an agent **418** that serves as a network proxy, where the agent **418** manages traffic to virtual IP addresses of a container system service. In this example, routing may be performed by the agent **418** and the container system service, which is user space.

[0269] The local area network **504** may be implemented with a variety of fabrics, devices, and protocols. The local area network **504** fabric may include Ethernet (802.3), wireless (802.11), and so forth. Data communication protocols for use over local area network **504** may include Transmission Control Protocol (TCP), User Datagram Protocol (UDP), Internet Protocol (IP), Hypertext Transfer Protocol (HTTP), Wireless Access Protocol (WAP), Handheld Device Transport Protocol (HDTP), Session Initiation Protocol (SIP), Real Time Protocol (RTP), among others.

[0270] As mentioned, the storage path routing module **500** may be configured to dynamically and intelligently route storage operations over network paths of different types based on one or more conditions of the container system **400**. In some implementations, the storage path routing module **500** may determine the one or more conditions based on network telemetry. Network telemetry may comprise data that is collected from one or more of: one or more computing resources within the container system **400**, one or more storage resources within the storage pool **406**, or one or more network resources associated with fabrics or devices among the computing resources and the storage resources.

[0271] For example, an agent **418** may periodically, or in response to an event, determine network telemetry data that is used by the controller **412** or by the container storage system **402**. A respective controller on a respective storage resource **424** may periodically, or in response to an event, determine the network telemetry data. Similarly, a device associated with a network resource may periodically, or in

response to an event, determine the network telemetry data. An event may be an indication of a change in status, such as a change into an error state, a reboot, an update, a change into a stable or normal state, or some other type of status change. Network telemetry data may also be determined by the container storage system **402** polling or pinging a node **420**, a storage resource **424**, or a network resource.

[0272] Network telemetry data may indicate, for individual network devices, status data. Status data may indicate a normal operating state, an error state, an average network traffic load, a current traffic load, a maximum traffic load, among other indications of an operating state. Network telemetry data may also indicate individual flow logs for individual network devices, where a flow log may indicate network traffic through the individual network device over a period of time. In a first example, for a given network path among a plurality of network paths, if each network device along the given network path is in a normal operating state, then the given network path is determined to be stable. In a second example, for a given network path among a plurality of network paths, if each network device along the given network path is both in normal operating state and has a traffic load below a threshold percentage of a maximum traffic load, then the given network path is determined to be stable.

[0273] The storage path routing module **500** may be configured to dynamically and intelligently route storage operations over network paths of different types based on any suitable set of one or more conditions of the container system **400**. In some embodiments, the conditions of the container system **400** may include network conditions of the storage area network **502** and/or the local area network **504**, such as latencies, throughput or data flow, etc. of the networks and/or network paths. In some embodiments, the conditions of the container system may include configuration settings of the container system **400**, containerized applications in the container system **400**, and/or container storage system **402**. For example, a configuration setting such as a defined number of replicas of a virtual volume may be considered by the storage path routing module **500** to select a storage path for a storage operation related to the virtual volume. Use of additional or alternative conditions of the container system **400** for intelligent routing of storage operations on network paths of different types are also contemplated.

[0274] In some implementations, the container storage system **402** may be configured to replicate data being stored within the storage pool **406**. For example, when a virtual volume **422** is provisioned, configuration information may indicate one or more of: a number of replicas, desired pool categories, storage pool selectors, or topology constraints. Based on the configuration information, storage provisioning by the container storage system **402** may include one or more of: distributing replicas across failure domains, distributing volumes across different nodes of a cluster to limit overprovisioning, and provisioning volumes among a specified type of storage. As a result of provisioning, the container storage system **402** determines a replication set of storage resources that host the virtual volume **422** mounted by containerized application such as containerized application **508**.

[0275] At a first time, given a virtual volume **422** mounted on a node **420**, a containerized application **508** may issue a first storage operation directed to the virtual volume **422**. For

example, the first storage operation may be a read operation, a write operation, an erase operation, or some type of I/O operation. The virtual volume **422** may be used by the containerized application **508** to store a dataset **510**. In this example, the dataset **510** is replicated by the container storage system **402** on multiple storage resources **424-1**, **424-2**, **424-M** within the storage pool **406**, as shown in FIG. 5.

[0276] The storage path routing module **500** may, in response to receiving the first storage operation and based on one or more first conditions, determine a network path over which to route the first storage operation. In some implementations, if the one or more first conditions are indicative of the first network path through the storage area network **502** being stable, then the storage path routing module **500** may determine to route the first storage operation over the first network path because the storage area network **502** is generally faster than the local area network **504**.

[0277] If the storage path routing module selects the first network path, then the storage path routing module **500** routes, over the first network path, the first storage operation to a first storage resource associated with the virtual volume **422**. The storage path routing module **500** may use a storage area network protocol that is used by a storage resource **424**.

[0278] In this example in which the dataset **510** on the virtual volume **422** is replicated, the storage path routing module **500** also routes the first storage operation to each replica storage resource **424**. If network paths through the storage area network **502** to each replica storage resources **424** are available and stable, the storage path routing module **500** routes the first storage operation to each replica storage resource **424** on network paths through the storage area network **502** and using a respective network protocol used by each respective storage resource **424**. In this example, the first storage operation is routed to the one or more storage resources **424** associated with the virtual volume **422** at the kernel level without the first storage operation being routed via the user space of the container system **400**.

[0279] In some implementations, where a dataset on a virtual volume is replicated among multiple host storage resources, the storage path routing module **500** may select multiple network paths. For example, based on the first one or more conditions, the storage path routing module **500** may determine to route the first storage operation over the storage area network **502** to a host storage resource **424-2**. Further, based on the first one or more conditions, the storage path routing module **500** may determine to route the first storage operation over the local area network **504** to another host storage resource **424-M**.

[0280] In some implementations, the storage path routing module **500** may switch a network path used to access a particular storage resource. For example, based on the first one or more conditions, the storage path routing module **500** may determine to use a first network path to access a first storage resource at a first time. However, at a subsequent time from the first time, based on a one or more changes to the first one or more conditions, the storage path routing module **500** may switch to route a subsequent storage operation over a second network path. A change in conditions may include a change to any of the conditions described above, including conditions that are determined based on network telemetry data, event data associated with network devices, or conditions of a network path.

[0281] In some implementations, network telemetry data may be used to train a machine learning model, where the model may be used to determine and/or predict one or more network conditions. For example, during operation, the storage path routing module 500 may use the model to predict whether a first network path or a second network path may fail or slow down. Based on the predictions for the first network path and the second network path, the storage path routing module 500 may select a network path for a storage operation instead of based solely on historical conditions or current conditions.

[0282] Continuing this example, at a second time, the containerized application 508 may issue a second storage operation directed to the virtual volume 422. The storage path routing module 500 may, in response to receiving the second storage operation, and based on one or more second conditions, determine a network path over which to route the second storage operation.

[0283] In some embodiments, the one or more second conditions may be determined based on updated network telemetry that is collected from one or more nodes and/or one or more storage resources. The network telemetry may be indicative of a status or health of a node, storage resource, or network component or path. In some examples, the one or more second conditions may be indicative of an error condition associated with the storage area network 502 and based on the one or more second conditions, the storage path routing module 500 selects the second network path over the local area network 504.

[0284] In some examples, the storage path routing module 500 may route the second storage operation to the first storage resource 424. Further, where the dataset 510 on the virtual volume 422 is replicated, the storage path routing module 500 also routes the second storage operation to each replica storage resource 424 according to a respective network protocol used by each respective storage resource 424. In an example in which the second conditions indicate an error condition associated with the storage area network 502, the second storage operation is routed to the one or more storage resources 424 associated with the virtual volume 422 at the user space level of the container system 400.

[0285] FIG. 6 depicts an example method 600 of storage path routing in a container system, in accordance with some embodiments of the present disclosure.

[0286] As described with respect to FIG. 5, a storage path routing module 500 of a container storage system 402 may dynamically route storage operations over a plurality of network paths within a container system 400. For example, a containerized application may access storage over a plurality of network paths, where each network path of the plurality of network paths may be distinct from another network path among the plurality of network paths. The plurality of network paths may be of different types and may include one or more paths over a local area network, one or more paths over a storage area network, and/or one or more paths over some other type of network.

[0287] The example method 600 includes, at 602, determining, within a container system, a storage operation directed to a storage resource. In some embodiments, the container system may be container system 400 described with respect to FIGS. 4A, 4B, and 5. A storage resource may be a storage resource within the storage pool 406, described with respect to FIGS. 4A, 4B, and 5. The container storage system 402 may respond to a request from a containerized

application to provision a volume of storage. In response, the container storage system 402 may provision storage from one or more storage resources 424 within the storage pool 406 and attach a block device on a node 420. The block device serves as a virtual volume for the containerized application. After being mounted, the virtual volume may be a target for storage operations issued by the containerized application.

[0288] In some examples, the virtual volume may be associated with a single storage resource and a dataset being used by a containerized application is stored on the single storage resource. In other examples, if a dataset being used by the containerized application is being replicated, then the virtual volume may be associated with a set of storage resources, where each of the storage resources in the set of storage resources stores a replica of the dataset.

[0289] The example method 600 further includes, at 604, determining one or more conditions associated with a first network path and with a second network path. The first network path may be over a storage area network, such as the storage area network 502. The second path may be over a local area network, such as the local area network 504. In some implementations, the one or more conditions used by the storage path routing module 500 to determine a network path for routing a storage operation may include any suitable combination (e.g., weighted combination) of conditions of the container system 400.

[0290] The example method 600 further includes, at 606, selecting, based on the one or more conditions, the first network path or the second network path. In some implementations, if the one or more conditions are indicative of the first network path being available and in a stable state, the first network path may be selected.

[0291] The example method 600 further includes, at 608, determining whether the first network path or the second network path has been selected.

[0292] The example method 600 further includes, at 610, routing the storage operation to the storage resource using the first network path over the storage area network 502, if the first network path was selected at 606. The storage path routing module 500 may route the storage operation over the storage area network 502 by communicating with a storage resource using a storage area network protocol.

[0293] The example method 600 further includes, at 612, routing the storage operation to the storage resource using the second network path over the local area network 504, if the second network path was selected at 606. Communication over the local area network 504 may be performed via the user space of the container system 400. For example, a node may have an agent 418 that serves as a network proxy, where the agent 418 manages traffic to virtual IP addresses of a container system service. In this example, routing may be performed by the agent 418 and the container system service, which is user space.

[0294] FIG. 7 depicts an example method 700 of storage path routing in a container system, in accordance with some embodiments of the present disclosure.

[0295] The example method 700 includes, at 702, determining, within a container system, a storage operation directed to a storage resource. Determining a storage operation directed to a storage resource may be performed similarly to determining, at 602 of FIG. 6, a storage operation directed to a storage resource.

[0296] The example method 700 further includes, at 704, determining one or more conditions associated with a plurality of network paths to the storage resource, where the plurality of network paths includes a first network path through a storage area network and a second network path through a container system communication network. Determining the one or more conditions associated with the plurality of network paths to the storage resource may be performed similarly to determining, at 604 of FIG. 6, one or more conditions associated with a first network path and a second network path.

[0297] The example method 700 further includes, at 706, routing based on the one or more conditions, the storage operation to the storage resource using the first network path. Routing the storage operation to the storage resource may be performed similarly to routing, at 610 or 612 of FIG. 6, a storage operation to a storage resource using either a first network path or a second network path.

[0298] FIG. 8 depicts an example method 800 of storage path routing in a container system, in accordance with some embodiments of the present disclosure.

[0299] In this example, a container storage system 402 may be deployed within the container system 400. The container storage system 402 may include a storage path routing module 500. The storage path routing module 500 may receive a storage operation directed to a virtual volume, and in response, the storage path routing module 500 routes the storage operation to host storage resource. The host storage resource may be a storage resource associated with the virtual volume.

[0300] The example method 800 includes, at 802, determining, by a container storage system, a request indicative of a storage operation directed to a storage resource. Determining a request indicative of a storage operation directed to a storage resource may be performed similarly to determining, at 602 of FIG. 6, a storage operation directed to a storage resource.

[0301] The example method 800 further includes, at 804, determining, by the container storage system, one or more conditions associated with a plurality of network paths to the storage resource, where the plurality of network paths includes a first network path through a storage area network and a second network path through a container system communication network. Determining the one or more conditions associated with the plurality of network paths to the storage resource may be performed similarly to determining, at 604 of FIG. 6, one or more conditions associated with a first network path and a second network path.

[0302] The example method 800 further includes, at 806, routing, by the container storage system and based at least on the one or more conditions, the storage operation to the storage resource using the first network path among the plurality of network paths. Routing the storage operation to the storage resource may be performed similarly to routing, at 610 or 612 of FIG. 6, a storage operation to a storage resource using either a first network path or a second network path.

[0303] In some embodiments, intelligent routing of storage operations over storage paths of different types in a container system, as described herein, may allow a user space of a software defined storage system in the container system, such as a container storage system running in the container system, to be bypassed by some storage operations. This allows a POD volume mounted on a node running

a containerized application to be directly connected to an actual target volume that is on the same node or remote from the node without going through the user space of the container system. When the target volume is on another node, SAN protocols may be used to set up IO paths to reach the target volume from the node where the POD volume is attached. In some examples, the user space of the container system may be involved only in setting up the path, and the kernel space can clone and route to the target volume now locally available via the in-kernel path. The in-kernel path may be faster and provide more bandwidth than a user space IO path. In some embodiments, the intelligent routing described herein may optimize the IO path to establish direct in-kernel IO cloning to all replicas with no further processing in user space. The intelligent routing described herein may be especially beneficial in a container system in which a container orchestrator moves PODs automatically and frequently.

[0304] In some embodiments, a virtual volume may go through a replication factor update. In response, the IO path may failover to a user space path and an in-kernel SAN path is not active. In this set of conditions, user space control agents (e.g., agents 418) manage bringing replicas into sync with each other. Once completed, an in-kernel SAN path may be reestablished.

[0305] In some embodiments, any failure in one of the replicated targets or a path to one of the replicated targets can cause failover of an in-kernel SAN path attached volume to a user space path. User space control agents (e.g., agents 418) may be configured to track conditions of volumes and once conditions of the volume in this example improve to a defined level, the IO path may be promoted back to the in-kernel SAN path to improve performance.

[0306] In some embodiments, intelligent routing of storage operations over storage paths of different types in a container system, as described herein, may be used for spread and/or proxy storage operations such as spread writes and proxy writes, for example. For a spread storage operation, a coordinator node may take full ownership in ensuring data is replicated across nodes. For a proxy storage operation, IO may be moved from the coordinator node to one of the designated target nodes having responsibility of replicating IO. A proxy virtual device exported over a storage area network may be set up by the control plane of container storage system 402 in some implementations. The proxy target exposes a proxy device that gets exported to the coordinator node such that all data movement in a proxy storage operation configuration is performed through a storage area network protocol. The proxy target node may be configured to use the storage area network 502 for data replication and/or other storage operations.

[0307] In some embodiments, in addition or alternative to intelligently routing storage operations over network paths of different types as described herein, one or more elements described herein (e.g., container storage system 402) may be configured to intelligently route storage operations locally on a node. For example, container storage system 402 may: receive a storage operation configured to be performed with respect to a storage resource; determine whether the storage resource is locally accessible at the same node as a virtual volume and/or a containerized application requesting or otherwise associated with the storage operation (e.g., determining whether the storage resource is mapped to the virtual volume); and route the storage operation based on this

determination such that, when the storage resource is locally accessible at the node, the routing is contained within the kernel space of the node (i.e., the route is direct and local to the storage resource without necessarily going over a network path).

[0308] As will be described and illustrated in more detail below with reference to FIGS. 9-13, the local routing of a storage operation may provide an additional storage path that may be selected, based on locality, and used when a target storage resource is located locally on the same node as a source of the storage operation. For example, the storage operation may be received from a container consuming virtual volume mounted on a node. Storage path routing module 500, or any other suitable element of a container system, may determine that a storage resource, such as a backing virtual volume, associated with the container consuming virtual volume is a target of the storage operation and is located on the same node. Based on this determination, the storage operation may be routed over a local, direct storage path, which may have low overhead and may avoid the storage operation being routed to user space where the storage operation would be rerouted over user space back to the same node.

[0309] In certain examples, one or more of the intelligent routing operations described herein may be performed in kernel space, within a kernel of an operating system of a node. For example, a storage system driver (or any other form of computer code) implemented in the kernel space of a node may be configured to intelligently route storage operations over any of a local, in-kernel storage path, an in-kernel SAN storage path to another node, and a container user space network path to another node based on one or more conditions, which may include any of the illustrative conditions described herein.

[0310] Advantages and features of the present disclosure can be further described by the following statements:

[0311] 1. A method comprising: receiving, by a container storage system executing on a node within a container system, a storage operation configured to be performed with respect to a storage resource; determining whether the storage resource is locally accessible at the node; and routing the storage operation based on the determining such that: when the storage resource is not locally accessible at the node, the routing is by way of both a kernel space and a user space of the node, and when the storage resource is locally accessible at the node, the routing is contained within the kernel space of the node.

[0312] 2. The method of statement 1, wherein: the container storage system includes a driver module executing in the kernel space of the node and a container module executing in the user space of the node; and the driver module receives the storage operation from a containerized application executing on the node.

[0313] 3. The method of either of statements 1 or 2, wherein the driver module of the container storage system implements a virtual block device accessible to the containerized application by way of a mount point exposed to the containerized application.

[0314] 4. The method of any of the preceding statements, wherein the container module of the container storage system implements a containerized storage application configured to execute on the node together with the containerized application that provides the storage operation.

[0315] 5. The method of any of the preceding statements, wherein the container module and the driver module of the container storage system are communicatively coupled by a Filesystem in Userspace (FUSE) interface executing in the kernel space of the node.

[0316] 6. The method of any of the preceding statements, wherein: the node is communicatively coupled with an additional node within the container system, the additional node executing an additional container storage system; the storage resource is locally accessible to the additional node rather than to the node; and the routing of the storage operation is by way of the kernel space of the node, the user space of the node, and both a user space and a kernel space of the additional node.

[0317] 7. The method of any of the preceding statements, wherein: the container storage system exposes a virtual block device executing in the kernel space of the node; the storage resource is locally accessible to the node and is managed by a file system executing within the kernel space of the node; and the virtual block device is configured to perform the storage operation by way of the file system.

[0318] 8. The method of any of the preceding statements, wherein: the container storage system exposes a virtual block device executing in the kernel space of the node; the storage resource is locally accessible to the node and is managed by a device mapper executing within the kernel space of the node; and the virtual block device is configured to perform the storage operation by way of the device mapper.

[0319] 9. The method of any of the preceding statements, wherein: the storage operation is further configured to be performed with respect to an additional storage resource on which a dataset stored in the storage resource is replicated; and when only the additional storage resource is locally accessible at the node, the routing of the storage operation includes routing the storage operation to both the storage resource and the additional storage resource by way of the kernel space and the user space of the node.

[0320] 10. The method of any of the preceding statements, wherein: the storage operation is further configured to be performed with respect to an additional storage resource on which a dataset stored in the storage resource is replicated; and when only the additional storage resource is locally accessible at the node, the routing of the storage operation includes containing a route of the storage operation to the additional storage resource within the kernel space of the node.

[0321] 11. The method of any of the preceding statements, wherein, when the routing is contained within the kernel space of the node, the routing uses a storage area network implemented on the node.

[0322] 12. The method of any of the preceding statements, wherein, when the routing is by way of both the kernel space and the user space of the node, the routing uses a local area network connecting the node to an additional node at which the storage resource is locally accessible.

[0323] 13. A container system comprising: a node that includes: a memory storing instructions, and one or more processors communicatively coupled to the memory and configured to execute the instructions to execute a container storage system; wherein the container storage system executing on the node is configured to perform a process comprising: receiving a storage operation configured to be performed with respect to a storage resource; determining

whether the storage resource is locally accessible at the node; and routing the storage operation based on the determining such that: when the storage resource is not locally accessible at the node, the routing is by way of both a kernel space and a user space of the node, and when the storage resource is locally accessible at the node, the routing is contained within the kernel space of the node.

[0324] 14. The container system of statement 13, wherein: the container storage system includes a driver module executing in the kernel space of the node and a container module executing in the user space of the node; and the driver module receives the storage operation from a containerized application executing on the node.

[0325] 15. The container system of either of statements 13 or 14, further comprising an additional node communicatively coupled with the node; wherein: the additional node is executing an additional container storage system; the storage resource is locally accessible to the additional node rather than to the node; and the routing of the storage operation is by way of the kernel space of the node, the user space of the node, and both a user space and a kernel space of the additional node.

[0326] 16. The container system of any of statements 13, 14, or 15, wherein: the container storage system exposes a virtual block device executing in the kernel space of the node; the storage resource is locally accessible to the node and is managed by a file system executing within the kernel space of the node; and the virtual block device is configured to perform the storage operation by way of the file system.

[0327] 17. The container system of any of statements 13, 14, 15, or 16, wherein: the container storage system exposes a virtual block device executing in the kernel space of the node; the storage resource is locally accessible to the node and is managed by a device mapper executing within the kernel space of the node; and the virtual block device is configured to perform the storage operation by way of the device mapper.

[0328] 18. The container system of any of statements 13, 14, 15, 16, or 17, wherein: the storage operation is further configured to be performed with respect to an additional storage resource on which a dataset stored in the storage resource is replicated; and when only the additional storage resource is locally accessible at the node, the routing of the storage operation includes routing the storage operation to both the storage resource and the additional storage resource by way of the kernel space and the user space of the node.

[0329] 19. The container system of any of statements 13, 14, 15, 16, 17, or 18, wherein: the storage operation is further configured to be performed with respect to an additional storage resource on which a dataset stored in the storage resource is replicated; and when only the additional storage resource is locally accessible at the node, the routing of the storage operation includes containing a route of the storage operation to the additional storage resource within the kernel space of the node.

[0330] 20. A non-transitory computer-readable medium storing instructions that, when executed, direct a processor of a node within a container system to perform a process comprising: receiving, by a container storage system executing on the node, a storage operation configured to be performed with respect to a storage resource; determining whether the storage resource is locally accessible at the node; and routing the storage operation based on the determining such that: when the storage resource is not locally

accessible at the node, the routing is by way of both a kernel space and a user space of the node, and when the storage resource is locally accessible at the node, the routing is contained within the kernel space of the node.

[0331] As mentioned above, a storage path routing module or other suitable driver module of a container storage system executing on a node (e.g., container storage system **402** executing on one of nodes **410** or **420**) may perform intelligent routing for storage operations directed to (i.e., configured to be performed with respect to) storage resources that are either locally accessible at the node (e.g., attached to the node and accessible through the kernel space of the node) or not locally accessible at the node (e.g., attached to a different node, etc.). Such storage operation routing may be implemented in any suitable manner to provide any of various valuable benefits, as will now be described in more detail with respect to FIGS. 9-13.

[0332] One issue that may be addressed by such routing is that when low-latency storage devices (e.g., devices using NVMe, etc.) are used as storage resources **424** in a storage pool **406**, a conventional transition of the storage operation from kernel space to user space may create excessive overhead that ultimately limits the total input/output operations per second (IOPS) that may be achieved. Accordingly, by bypassing the transitions from kernel space to user space and from user space back to kernel space and, instead, containing the routing of storage operations within the kernel space of a node where possible (e.g., in situations where the storage resource is locally accessible at the node), greater efficiency and throughput (e.g., greater IOPS) with less latency may be achieved as the storage operations are performed.

[0333] These benefits may arise for any container storage system that includes modules or components in both a kernel space (e.g., a driver module) and a user space (e.g., a containerized application). Additionally, certain storage systems implemented in user space and relying on Filesystem in Userspace (FUSE) drivers within the kernel space may similarly leverage benefits provided by the kernel-contained routing paths to locally accessible storage resources described herein. For example, the overhead eliminated by routing storage operations only within kernel space instead of sending them to be handled by a container module of the container storage system in user space may, in some examples, allow for at least double the throughput (e.g., from 100 kIOPs for routing through user space to >200 kIOPs for routing contained within kernel space) while also potentially improving latency, reliability, and so forth. At the same time, for storage operations configured to be performed with respect to storage resources that are not locally accessible at the node (or when other issues occur that prevent the kernel-only routing path from being used), conventional routing by way of the user space module may still be used as a default or fallback mode providing all the benefits for which it has conventionally been relied on.

[0334] FIG. 9 depicts an example method **900** of routing a storage operation in a container system in accordance with principles described herein. As shown, method **900** may be performed, in full or in part, by a container storage system executing on a node within a container system, such as a container storage system **402** executing on a node **410** or **420** within container system **400**, as illustrated and described above. While method **900** shows illustrative operations according to one implementation, other implementations

may omit, add to, reorder, and/or modify any of the operations shown in method **900**. In some examples, multiple operations shown in method **900** may be performed concurrently (e.g., in parallel) with one another, rather than being performed sequentially as illustrated and/or described. One or more of the operations shown in FIG. **9** may be performed by a storage path routing module or other driver module within the container storage system **402** (e.g., storage path routing module **500** or an implementation thereof). Each of operations **902-910** of method **900** will now be described.

[0335] At operation **902**, the container storage system **402** may receive a storage operation configured to be performed with respect to a storage resource. For example, the storage operation may originate at a containerized application within the node (e.g., one of containerized applications **404** or another suitable containerized application) or from any other suitable source. As has been mentioned (and as will be described in more detail below), the container storage system **402** may implement a virtual block device and/or other driver software executing within kernel space of the node, and the receiving of the storage operation at operation **902** may be performed by one of these components within kernel space. As such, it may be necessary for the storage operation received in kernel space to transition to user space (e.g., by way of a FUSE module or the like) if the storage operation is to be handled in user space by, for instance, a containerized module of the container storage system **402** that executes in user space. While this may be one way that storage operations are handled, various benefits described above may be gained when intelligent routing is applied to achieve a more efficient and direct routing of the storage operation in certain conditions.

[0336] To this end, at operation **904**, the container storage system **402** may determine whether the storage resource is locally accessible at the node. For instance, in the example of FIG. **4A**, it is shown that storage resource **424-1** is locally accessible to node **420-1**. Though other storage resources **424-2** through **424-M** within the storage pool **406** are also shown to be available to implement virtual volume **422-1** for node **420-1**, these other storage resources are shown to not be locally accessible to the node (e.g., though they may be locally accessible to other nodes **420** in the cluster **408**). Accordingly, at operation **904**, the determination is whether the storage resource to which the particular storage operation received at operation **902** is directed is locally accessible to the node (e.g., as storage resource **424-1** is to node **420-1**) or must be accessed more remotely (e.g., as would be the case for storage resources **424-2** through **424-M** in relation to node **420-1**).

[0337] At operation **906**, the container storage system **402** may route the storage operation based on the determination made at operation **904**. Specifically, the container storage system **402** may route the storage operation such that: 1) when the storage resource is not locally accessible at the node, a routing operation **908** is by way of both a kernel space and a user space of the node, and 2) when the storage resource is locally accessible at the node, a routing operation **910** is contained within the kernel space of the node. Accordingly, the container storage system **402** may often (e.g., when the storage resource is locally accessible at the node) avoid performing the overhead of transitioning the storage operation to user space to be handled by a containerized storage application or other container module of the container storage system **402** executing in user space. In

these situations, a faster and more efficient route that avoids that overhead may be used to achieve various benefits that have been described and will be made more apparent below.

[0338] While not explicitly shown in FIG. **9**, it will be understood that the intelligent storage operation routing implemented by method **900** may be enabled or disabled based on user commands (e.g., a command to enable or disable a “fast path” mode for storage operations, etc.), predesignated user preferences, system settings, system or network conditions, the nature of a given storage operation (e.g., how important it is for the storage operation to be performed with optimum latency and throughput, etc.), and/or any other suitable circumstances or conditions as may serve a particular implementation.

[0339] To illustrate a context in which method **900** may be performed, FIG. **10** shows a configuration **1000** in which various routes through a kernel space **1002** and/or a user space **1004** of a node **1006** may be implemented for routing a storage operation **1008** to a storage resource **424** in accordance with principles described herein. As shown in FIG. **10**, a storage path routing module **500** (e.g., an implementation of any routing module described herein) may be implemented within a container storage system **402** (e.g., an implementation of any of the container storage systems described herein) executing on the node **1006** (e.g., an implementation of any of nodes **410** or **420** or other computing nodes described herein). The container storage system **402** is shown to span both the kernel space **1002** and the user space **1004** of the node **1006**, indicating (as will be described in more detail below) that the container storage system **402** may include modules configured to execute in both of these spaces of node **1006**. As further shown, the storage path routing module **500** may at least partially execute within the kernel space **1002** and, within the kernel space, may receive a storage operation **1008** that is provided to container storage system **402**.

[0340] Upon receiving the storage operation **1008** (operation **902**), storage path routing module **500** may identify a storage resource to which the storage operation is directed (i.e., a storage resource with respect to which the storage operation is configured to be performed) and may determine whether that storage resource is locally accessible at the node **1006**. For purposes of this example (and other examples that will be described below), a local storage resource **424-L** (‘L’ signifying “local”) will be understood to represent a storage resource that is accessible at the particular node **1006**, while a non-local storage resource **424-R** (‘R’ signifying “remote”) will be understood to represent a storage resource that is not accessible at the particular node **1006** (e.g., though it may be locally accessible to a different node **1006**, as will be described and illustrated in an example below).

[0341] As such, FIG. **10** illustrates three different routes or paths by way of which storage operation **1008** may be routed to be performed with respect to whichever of these storage resources it is directed. As shown, these routes are depicted using dashed arrows that are distinguishable from solid arrows meant to represent communicative connections between different components in the figure. Each of these routes will now be described.

[0342] First, a route **1010** is depicted that represents a route used by certain container storage systems (e.g., container storage systems implemented as containerized applications) to route storage operations to a local storage

resource such as local storage resource **424-L**. As shown, route **1010** extends between storage path routing module **500** and local storage resource **424-L** by way of both the kernel space **1002** and the user space **1004**. For example, upon receiving storage operation **1008** and determining that the storage resource to which it is directed is local storage resource **424-L**, storage path routing module **500** may send data representative of storage operation **1008** from kernel space **1002** to user space **1004** to be handled there and, ultimately, to be performed with respect to local storage resource **424-L**. For example, in user space **1004**, a user space module of container storage system **402** may receive and route data representative of the storage operation **1008** back through the kernel space to the local storage resource **424-L**. In some implementations, route **1010** may serve as a useful default path, since a module of container storage system **402** executing in user space **1004** may be well-situated to handle both storage operations directed to local and non-local storage resources. However, as has been mentioned, a route through both the kernel space and the user space such as route **1010** may be associated with more overhead than may be necessary when the storage resource is accessible locally.

[0343] A second route, route **1012**, is depicted to represent how a container storage system (e.g., a container storage system implemented as a containerized application) may route storage operations to non-local storage resources such as remote storage resource **424-R**. Similar to route **1010**, route **1012** is shown to extend between storage path routing module **500** and remote storage resource **424-R** by way of both the kernel space **1002** and the user space **1004**. In this case, however, the overhead associated with transitioning a storage operation **1008** from kernel space **1002** to user space **1004** may be more justified, since remote storage resource **424-R** is not locally accessible at this particular node **1006** (e.g., is not accessible through kernel space **1002** alone).

[0344] Third, a route **1014** between storage path routing module **500** and local storage resource **424-L** is shown to be contained within kernel space **1002** of node **1006**. This route will be understood to represent a more optimal path (e.g., a “fast path”) that may be used when suitable conditions exist, such as when the storage resource to which storage operation **1008** is directed is locally accessible at node **1006** (e.g., local storage resource **424-L**) and when the communication path between the container storage system **402** and the local storage resource **424-L** (which may include a storage area network (SAN) or other such communicative components) is operational and available for use. As has been described, while either of routes **1010** or **1014** may allow a storage operation **1008** to successfully be routed to local storage resource **424-L**, route **1014** may allow container storage system **402** to avoid the overhead of transitioning the storage operation into and back out of user space **1004**, and may therefore provide certain throughput and efficiency benefits that have been described. It will be understood that, while route **1014** may be preferred when the conditions are such that it can be used, route **1010** may provide a fallback option when conditions change (e.g., when the “fast path” mode is disabled, etc.).

[0345] While the examples described thus far in FIGS. 9 and 10 have assumed that a given storage operation **1008** is configured to be performed with respect to a singular storage resource, it will be understood (as mentioned above) that this may not always be the case. Rather, in some examples, a

storage operation **1008** may be directed to a plurality of storage resources on which data is (or is intended to be) replicated. In other words, a storage operation **1008** may not only be configured to be performed with respect to a storage resource but may also be further configured to be performed with respect to an additional storage resource on which a dataset stored in the storage resource is replicated. In this scenario, the routing performed by container storage system **402** (and by storage path routing module **500** in particular) may use the same type of route or different types of routes (e.g., routes **1010**, **1012**, and/or **1014**) as may serve a particular implementation.

[0346] To illustrate, FIG. 11 shows example implementations for routing storage operations that are configured to be performed with respect to multiple storage resources. Specifically, a particular storage operation **1008** is shown to be directed to a plurality of different storage resources on which a dataset is replicated. Each of the storage resources to which storage operation **1008** is directed is shown to be represented by a different box labeled “To Storage Resource” (two of which are explicitly shown in FIG. 11, though the ellipsis in storage operation **1008** indicates that more than two may be included in other examples) within storage operation **1008**. Three possible scenarios **1102**, **1104**, and **1106** are shown in a table in which the “To Storage Resource” boxes are either in a column **1108** (“Locally Accessible”) or a column **1110** (“Remotely Accessible”). In scenario **1102**, storage operation **1008** is shown to be configured to be performed with respect to storage resources that are both locally accessible (such as storage resource **424-L**). In scenario **1104**, storage operation **1008** is shown to be configured to be performed with respect to storage resources that are both remotely accessible (such as storage resource **424-R**). In scenario **1106**, storage operation **1008** is shown to be configured to be performed with respect to one storage resource that is locally accessible (in column **1108**) and one storage resource that is remotely accessible (in column **1110**).

[0347] On the right side of the table, two potential implementations **1112** and **1114** are shown for how to handle a multi-storage-resource storage operation such as this example of storage operation **1008**. Implementation **1112** is labeled as a “Different Route Implementation,” indicating that, when applicable, different types of routes (selected from routes **1010**, **1012**, and **1014**) may be employed to route the storage operation **1008** to storage resources accessible in different ways. In contrast, implementation **1114** is labeled as a “Same Route Implementation,” indicating that the same type of route (selected from routes **1010**, **1012**, and **1014**) may be employed to route the storage operation **1008**, even if the storage resources are accessible in different ways.

[0348] When a storage operation **1008** is configured to be performed with respect to both a storage resource and an additional storage resource (on which a dataset stored in the storage resource is replicated), and when only the additional storage resource is locally accessible at the node (i.e., in scenario **1106** when the storage resources are accessible in different ways), the different route implementation **1112** may use the routing of method **900** (by way of the kernel space and the user space) to route the storage operation **1008** to the storage resource (i.e., the remote storage resource in this example) while also containing a route of the storage

operation to the additional storage resource (i.e., the local storage resource in this example) within the kernel space of the node.

[0349] In contrast, the same route implementation 1114 for this same scenario (in which the storage operation 1008 is configured to be performed with respect to both the storage resource and the additional storage resource on which a dataset stored in the storage resource is replicated, and in which only the additional storage resource is locally accessible at the node) may operate differently. For example, the same route implementation 1114 may route the storage operation to both the storage resource (i.e., the remote storage resource in this example) and the additional storage resource (i.e., the local storage resource in this example) by way of the kernel space and the user space of the node.

[0350] To illustrate, FIG. 11 shows that the same types of routes are used for both implementations 1112 and 1114 in scenario 1102 and in 1104, while these implementations differ in how they handle scenario 1106. Specifically, as shown in scenario 1102 (when both storage resources are locally accessible), both implementations 1112 and 1114 contain the routing to the kernel space 1002 (“Kernel Space Only (e.g., Route 1014)”). Similarly, as shown in scenario 1104 (when both storage resources are remotely accessible), both implementations 1112 and 1114 route the storage operation by way of both the kernel space 1002 and the user space 1004 (“Kernel Space/User Space (e.g., Route 1012)”). When the replication occurs both locally and remotely as in scenario 1106, however, FIG. 11 shows that implementations 1112 and 1114 differ in their approach. In the different route implementation 1112, the locally accessible storage resource route is contained to the kernel space 1002 (“Kernel Space Only (e.g., Route 1014)”) while the remotely accessible storage resource route is by way of both the kernel space 1002 and the user space 1004 (“Kernel Space/User Space (e.g., Route 1012)”). In contrast, in the same route implementation 1114, both of the routes are shown to be by way of the kernel space 1002 and the user space 1004 (“Kernel Space/User Space (e.g., Route 1012)”), even though one of the storage resources is locally accessible.

[0351] The method 900 described with respect to FIG. 9 and the various aspects described with respect to FIGS. 10 and 11 may be implemented in any manner as may serve a particular implementation. To illustrate additional detail by way of which configuration 1000 may be implemented in certain embodiments, FIG. 12 shows another example configuration 1200 in which the same routes 1010, 1012, and 1014 through the kernel space 1002 and/or the user space 1004 of the node 1006 may be implemented for routing the storage operation 1008 to either or both of storage resources 424-L and/or 424-R.

[0352] As shown within a dashed outline of the container storage system 402 in FIG. 12, the container storage system 402 may include various components and/or modules executing in both the kernel space 1002 and the user space 1004 of the node 1006. For example, FIG. 12 shows a driver module 1202 executing in the kernel space 1002 of the node and a container module 1204 executing in the user space 1004 of the node.

[0353] Driver module 1202 may implement the storage path routing module 500 described above (e.g., as shown in FIG. 10), and, as such, may be configured to implement method 900 and other routing-related operations described herein. Driver module 1202 may also expose a virtual block

device and one or more storage management modules, as will be described in more detail below.

[0354] Container module 1204 may implement (e.g., may be implemented as) a containerized storage application configured to execute on the node 1006 together with one or more other containerized applications. For example, container module 1204 may execute in a container runtime alongside any of the other containerized applications described herein, including a containerized application 404 that provides the storage operation 1008 (described in more detail below). In such examples, container module 1204 may execute on any computing hardware and may be managed by any suitable container orchestrator or container orchestration platform (e.g., Kubernetes, etc.). In other examples, container module 1204 may not be containerized but, rather, may be implemented as a standalone application executing within the user space 1004 of the node 1006.

[0355] As further depicted in the implementation of container storage system 402 in FIG. 12, the container module 1204 and the driver module 1202 of the container storage system 402 may be communicatively coupled by a Filesystem in Userspace (FUSE) interface executing in the kernel space 1002 of the node 1006. The FUSE interface is labeled FUSE 1206 in FIG. 12 and may serve as a software interface for the operating system to allow containerized applications and/or other applications executing in the user space 1004 to implement filesystem utilities without having write access to kernel code. In this way, FUSE 1206 may provide a kernel interface for filesystem code executing in the user space 1004 (e.g., container module 1204).

[0356] While FIG. 10 did not explicitly show a source that provides the received storage operation 1008, FIG. 12 illustrates that the driver module 1202 may receive the storage operation 1008 from a containerized application 404 that is executing on the node 1006. For instance, as shown, the storage operation 1008 may be directed to driver module 1202 by way of a mount point 1208 that is provided to the containerized application 404 to allow the containerized application to read and write data to persistent storage (e.g., to virtual volumes implemented by storage resources 424) in any of the ways described herein.

[0357] As shown, mount point 1208 may be configured to allow data such as the storage operation 1008 to be passed to a virtual block device 1210 provided by the driver module 1202. For example, the driver module 1202 of the container storage system 402 may implement the virtual block device 1210 so as to be accessible to the containerized application 404 by way of the mount point 1208 exposed to the containerized application 404. The virtual block device 1210 may serve as a virtual file configured to provide buffered access to a virtual volume accessible by way of the container storage system 402 (e.g., a virtual volume implemented by one or more local and/or remote storage resources 424). The virtual block device may be configured to allow read/write access to a storage block of any size and/or alignment.

[0358] A storage management utility 1212 is also shown in FIG. 12 to be further implemented within the kernel space 1002 by the driver module 1202. The storage management utility 1212 may be configured to provide an interface between driver module 1202 and other utilities of the node executing in the kernel space 1002 (e.g., built-in operating system utilities, driver code from other driver modules, etc.) in any suitable way. For instance, in certain examples, the storage management utility 1212 may interface with oper-

ating system utilities used by the driver module **1202** (e.g., established filesystems, device mappers, etc.), while, in other examples, the storage management utility **1212** may interface with or fully or partially implement code customized for the container storage system **402** to allow the storage management utility **1212** to interface with the container module **1204** in unique ways different from established utilities.

[0359] In some implementations, storage management utility **1212** may be configured to interface with (or may include or implement) a file system (e.g., an ext4 file system, an xfs file system, a btrfs file system, etc.) executing within the kernel space **1002**. For example, in implementations in which the container storage system **402** exposes the virtual block device **1210** executing in the kernel space **1002**, the storage resource **424-L** locally accessible to the node **1006** may be managed by one or more file systems and the virtual block device **1210** may be configured to perform the storage operation **1008** by way of these one or more file systems.

[0360] Additionally or alternatively, storage management utility **1212** may be configured to interface with a device mapper executing within the kernel space **1002** of the node (e.g., a thin-provisioning device mapper such as a dmthin utility, etc.). For example, in implementations in which the container storage system **402** exposes the virtual block device **1210** executing in the kernel space **1002**, the storage resource **424-L** locally accessible to the node **1006** may be virtualized using a device mapper to create a virtual storage pool that functions as a backing datastore that is exposed to a container system. The container storage system **402** may thinly provision logical volumes in the datastore, which logical volumes may function as backing volumes for virtual volumes that are provided by the container storage system **402** to the container system for use by containerized applications running in the container system. The virtual storage pool and the volumes created in the virtual storage pool may be in device namespace of the node. The virtual block device **1210** may be configured to route the storage operation **1008** to an appropriate logical volume in the virtualized storage pool, without going through a filesystem namespace. Thus, storage operations received on virtual block device **1210** may be provided directly to the device-mapped block device to be processed entirely within kernel space **1002**.

[0361] Along with file systems and device mappers as have been described, storage management utility **1212** may also be configured to interface with (or include or implement at least portions of) other utilities such as logical volume manager utilities (e.g., lvm), RAID utilities (mdraid), and/or any other suitable software as may serve a particular implementation.

[0362] While non-local storage resources such as remote storage resource **424-R** are not locally accessible to a container storage system **402** at one particular node **1006**, it will be understood that these storage resources may, in certain cases, be locally accessible to another container storage system **402** at another node in the same cluster (e.g., cluster **408**) and/or within the same container system (e.g., container system **400**). Accordingly, when a route such as route **1012** is selected for routing a storage operation, the routing may traverse not only the kernel space and user space of the particular node **1006**, but also the kernel space and user space of the additional node at which the storage resource is locally accessible.

[0363] To illustrate, FIG. **13** shows another configuration **1300** in which the same routes **1010**, **1012**, and **1014** that have been described are implemented. As in configurations **1000** and **1200**, configuration **1300** shows that route **1010** is by way of both the kernel space **1002** and/or the user space **1004** of the node **1006** and that route **1014** is contained within the kernel space **1002**. Configuration **1300** shows additional detail, however for how route **1012** may be in communication with remote storage resource **424-R**, as well as how routes **1010** and **1014** may be in communication with local storage resource **424-L**.

[0364] More specifically, as shown in configuration **1300**, the node **1006** referenced in configurations **1000** and **1200** is relabeled as node **1006-1** and is shown to be communicatively coupled with an additional node labeled as node **1006-2** within the container system. The additional node **1006-2** is shown to be executing an additional container storage system **402** that includes analogous modules and components as the container storage system **402** of node **1006-1**, including, for instance, a driver module **1202**, a container module **1204**, and a FUSE interface **1206**. To differentiate the components of node **1006-1** from analogous components of node **1006-2**, '-1' and '-2' postfixes are employed in like manner as with the nodes **1006**. That is, while node **1006-1** is shown to include a container storage system **402-1** with a driver module **1202-1**, a container module **1204-1**, and a FUSE interface **1206-1**, node **1006-2** is shown to include a container storage system **402-2** with a driver module **1202-2**, a container module **1204-2**, and a FUSE interface **1206-2**.

[0365] While remote storage resource **424-R** is not locally accessible to node **1006-1**, FIG. **13** shows that remote storage resource **424-R** is locally accessible to the additional node **1006-2** in an analogous way as local storage resource **424-L** is locally accessible to node **1006-1**. Specifically, as shown, each of these storage resources may be accessed by way of a respective storage area network (SAN) **1302** (i.e., SAN **1302-1** for node **1006-1** and SAN **1302-2** for node **1006-2**). In this way, when the routing is contained within the kernel space **1002** of node **1006-1** (e.g., using route **1014**, as has been described), the routing may use the storage area network **1302-1** implemented on the node **1006-1**. On the other hand, when the routing is to remote storage resource **424-R** attached to node **1006-2** (such as is the case for route **1012**), the routing of the storage operation **1008** may be by way of the kernel space **1002** of the node **1006-1**, the user space **1004** of the node **1006-1**, and both a user space and a kernel space of the additional node **1006-2**.

[0366] As further shown in configuration **1300** by the example of route **1012**, when the routing is by way of both the kernel space **1002** and the user space **1004** of the node **1006-1**, the routing may use a local area network (LAN) **1304** connecting the node **1006-1** to the additional node **1006-2** at which the remote storage resource **424-R** is locally accessible. LAN **1304** may be implemented in any suitable way and involve any suitable messaging protocols or brokers, including protocols that do not use a dedicated message broker (e.g., Zero Message Queue (ZMQ), etc.).

[0367] In the preceding description, various exemplary embodiments have been described with reference to the accompanying drawings. It will, however, be evident that various modifications and changes may be made thereto, and additional embodiments may be implemented, without departing from the scope of the invention as set forth in the

claims that follow. For example, certain features of one embodiment described herein may be combined with or substituted for features of another embodiment described herein. The description and drawings are accordingly to be regarded in an illustrative rather than a restrictive sense.

What is claimed is:

1. A method comprising:

receiving, by a container storage system executing on a node within a container system, a storage operation configured to be performed with respect to a storage resource;

determining whether the storage resource is locally accessible at the node; and

routing the storage operation based on the determining such that:

when the storage resource is not locally accessible at the node, the routing is by way of both a kernel space and a user space of the node, and

when the storage resource is locally accessible at the node, the routing is contained within the kernel space of the node.

2. The method of claim 1, wherein:

the container storage system includes a driver module executing in the kernel space of the node and a container module executing in the user space of the node; and

the driver module receives the storage operation from a containerized application executing on the node.

3. The method of claim 2, wherein the driver module of the container storage system implements a virtual block device accessible to the containerized application by way of a mount point exposed to the containerized application.

4. The method of claim 2, wherein the container module of the container storage system implements a containerized storage application configured to execute on the node together with the containerized application that provides the storage operation.

5. The method of claim 2, wherein the container module and the driver module of the container storage system are communicatively coupled by a Filesystem in Userspace (FUSE) interface executing in the kernel space of the node.

6. The method of claim 1, wherein:

the node is communicatively coupled with an additional node within the container system, the additional node executing an additional container storage system;

the storage resource is locally accessible to the additional node rather than to the node; and

the routing of the storage operation is by way of the kernel space of the node, the user space of the node, and both a user space and a kernel space of the additional node.

7. The method of claim 1, wherein:

the container storage system exposes a virtual block device executing in the kernel space of the node;

the storage resource is locally accessible to the node and is managed by a file system executing within the kernel space of the node; and

the virtual block device is configured to perform the storage operation by way of the file system.

8. The method of claim 1, wherein:

the container storage system exposes a virtual block device executing in the kernel space of the node;

the storage resource is locally accessible to the node and is virtualized by a device mapper executing within the kernel space of the node; and

the virtual block device is configured to route the storage operation to the virtualized storage resource.

9. The method of claim 1, wherein:

the storage operation is further configured to be performed with respect to an additional storage resource on which a dataset stored in the storage resource is replicated; and

when only the additional storage resource is locally accessible at the node, the routing of the storage operation includes routing the storage operation to both the storage resource and the additional storage resource by way of the kernel space and the user space of the node.

10. The method of claim 1, wherein:

the storage operation is further configured to be performed with respect to an additional storage resource on which a dataset stored in the storage resource is replicated; and

when only the additional storage resource is locally accessible at the node, the routing of the storage operation includes containing a route of the storage operation to the additional storage resource within the kernel space of the node.

11. The method of claim 1, wherein, when the routing is contained within the kernel space of the node, the routing uses a storage area network implemented on the node.

12. The method of claim 1, wherein, when the routing is by way of both the kernel space and the user space of the node, the routing uses a local area network connecting the node to an additional node at which the storage resource is locally accessible.

13. A container system comprising:

a node that includes:

a memory storing instructions, and

one or more processors communicatively coupled to the memory and configured to execute the instructions to execute a container storage system;

wherein the container storage system executing on the node is configured to perform a process comprising:

receiving a storage operation configured to be performed with respect to a storage resource;

determining whether the storage resource is locally accessible at the node; and

routing the storage operation based on the determining such that:

when the storage resource is not locally accessible at the node, the routing is by way of both a kernel space and a user space of the node, and

when the storage resource is locally accessible at the node, the routing is contained within the kernel space of the node.

14. The container system of claim 13, wherein:

the container storage system includes a driver module executing in the kernel space of the node and a container module executing in the user space of the node; and

the driver module receives the storage operation from a containerized application executing on the node.

15. The container system of claim 13, further comprising an additional node communicatively coupled with the node; wherein:

the additional node is executing an additional container storage system;

the storage resource is locally accessible to the additional node rather than to the node; and
 the routing of the storage operation is by way of the kernel space of the node, the user space of the node, and both a user space and a kernel space of the additional node.

16. The container system of claim **13**, wherein:
 the container storage system exposes a virtual block device executing in the kernel space of the node;
 the storage resource is locally accessible to the node and is managed by a file system executing within the kernel space of the node; and
 the virtual block device is configured to perform the storage operation by way of the file system.

17. The container system of claim **13**, wherein:
 the container storage system exposes a virtual block device executing in the kernel space of the node;
 the storage resource is locally accessible to the node and is virtualized by a device mapper executing within the kernel space of the node; and
 the virtual block device is configured to route the storage operation to the virtualized storage resource.

18. The container system of claim **13**, wherein:
 the storage operation is further configured to be performed with respect to an additional storage resource on which a dataset stored in the storage resource is replicated; and
 when only the additional storage resource is locally accessible at the node, the routing of the storage operation includes routing the storage operation to both

the storage resource and the additional storage resource by way of the kernel space and the user space of the node.

19. The container system of claim **13**, wherein:
 the storage operation is further configured to be performed with respect to an additional storage resource on which a dataset stored in the storage resource is replicated; and

when only the additional storage resource is locally accessible at the node, the routing of the storage operation includes containing a route of the storage operation to the additional storage resource within the kernel space of the node.

20. A non-transitory computer-readable medium storing instructions that, when executed, direct a processor of a node within a container system to perform a process comprising:
 receiving, by a container storage system executing on the node, a storage operation configured to be performed with respect to a storage resource;
 determining whether the storage resource is locally accessible at the node; and
 routing the storage operation based on the determining such that:

when the storage resource is not locally accessible at the node, the routing is by way of both a kernel space and a user space of the node, and

when the storage resource is locally accessible at the node, the routing is contained within the kernel space of the node.

* * * * *